DOCUMENT RESUME

ED 166 223 TM 008 172

AUTHOR Davis, Charles E.; And Others

TITLE A Primer on Decision Analysis for Individually Prescribed Instruction. ACT Technical Bulletin No. 17.

INSTITUTION American Coll. Testing Program, Iowa City, Iowa. Research and Development Div.

SPONS AGENCY Office of Education (DHEW), Washington, D.C.
PUB DATE Oct 73
GRANT OEG-0-72-0711
NOTE 104p.

EDRS PRICE MF-$0.83 HC-$6.01 Plus Postage.
DESCRIPTORS *Bayesian Statistics; *Cutting Scores; *Decision Making; Logic; Mastery Tests; *Mathematical Models; Test Interpretation
IDENTIFIERS Decision Theoretic Models; Utility Theory

ABSTRACT

A coherent system of decision making is described that may be incorporated into an instructional sequence to provide a supplement to the experience-based judgment of the classroom teacher. The elements of this decision process incorporate prior information such as a teacher's past experience, experimental results such as a test score, and preferences or values. Beginning with a basic example in which the prior information is the probability that a student having taken a course will have mastered it; the experimental result is a test item score, and the values expressed are the costs of wrong pass-fail decisions. Equations and numerical examples demonstrate the solution to this and to more complex problems. Following these examples is another group of examples that abandons the notion of loss or costs in favor of the notion of the utility functions for various decisions. There are discussions of the formal structure of decision theory and of utility theory. Many of the examples are complex and much of the argument is expressed in integrals or summation equations. Topics include: normal form analysis and extensive form analysis; continuous posterior probabilities; a binomial model; threshold utility; linear utility; quadratic, exponential, and squared exponential utility; and three action examples. (CTM)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED166223

ACT TECHNICAL BULLETIN NO. 17

TM

U.S. DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

A Primer on Decision Analysis for Individually Prescribed Instruction

by

Charles E. Davis
The University of Iowa

James Hickman
The University of Wisconsin

and

Melvin R. Novick
The American College Testing Program
and
The University of Iowa

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
R. Ferguson
ACT
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) AND USERS OF THE ERIC SYSTEM."

The Research and Development Division
The American College Testing Program
P. O. Box 168, Iowa City, Iowa 52240

October, 1973

IM008 172

A Primer on Decision Analysis for Individually Prescribed Instruction[1]

Charles E. Davis
The University of Iowa

James Hickman
The University of Wisconsin

and

Melvin R. Novick

The American College Testing Program
and
The University of Iowa

1. Introduction

If you are lucky, and certain that your luck will hold, you should read no further. Our subject is decision making and those who are inherently lucky will have no need to attempt to take advantage of a logical system for decision making. However, we feel that logical thought has been successful in so many areas of human activity that it deserves a chance even in those areas where arbitrary rules or intuition seem to have prevailed, and perhaps been moderately successful. In particular, we feel that the implementation of Individually Prescribed Instruction can be given a greater payoff if a coherent system of decision making is incorporated into the instructional sequence to provide a supplement to the experience-honed judgment of the classroom teacher.

In the currently popular language of systems engineering, the decision making process might be viewed as a black box. The black box contains an input hole for prior information about the environment in which the decision will be made and evaluated, a second hole for new experimental results designed

[1]The research reported herein was performed pursuant to Grant No. OEG-0-72-0711 with the Office of Education, U. S. Department of Health, Education, and Welfare, Melvin R. Novick, Principal Investigator. Contractors undertaking such projects under Government sponsorship are encouraged to express freely their professional judgment in the conduct of the project. Points of view or opinions stated do not, therefore, necessarily represent official Office of Education position or policy. We are grateful to Nancy S. Peterson for carefully reading and correcting this manuscript.

explicitly to cast more light on the uncertain environment, and finally a hole for the decision maker's preferences for the possible consequences of decisions. The output of this black box is a decision or action that will ultimately lead to a distribution of consequences that will, hopefully, be high on the decision maker's preference scale. Our business will be to get enough of a look at what goes on inside the black box so that we can see how the teacher and student can use prior information, experimental results, and preferences in a reasonable way to generate decisions having favorable consequences.

Building an all-purpose black box for decision making has been a major project in recent years. Contributors to the project have included economists, mathematicians, philosophers, psychologists, and statisticians. It would take a major treatise to adequately trace this project. We must, therefore, limit our survey to the major references in psychology and education relevant to our task. Those elements of decision analysis which require that existing information about the process under study be quantified in order that new experimental information may be coherently combined with the existing information, have already been persuasively presented in the literature of psychology and education [Edwards, Lindman, and Savage (1963)]. However, the requirement that preferences for various possible distributions of consequences be formulated coherently and expressed as a numerical valued utility function, has not been emphasized in the literature of education with the singular exception of the prophetic text, Psychological Tests and Personnel Decisions, by Cronbach and Gleser (1957, 1965).

Many decision makers in education may feel that business managers, with the market at hand to evaluate the outcomes associated with their decisions, are in a more favorable position to make value comparisons among the

distributions of consequences than they are. For example, the values attached to the consequences of a decision to implement a specific education program, or to advance or retain a particular student at a given level, seem to them to defy simple quantification. Such decisions appear to have so many possible ramifications that to formulate preferences with any degree of consistency is simply impossible.

Nevertheless, it is our position that most data-collecting activities in education are designed to influence decisions. In fact, to simply collect data without the objective of ultimately modifying a course of action, would seem wasteful. We further assert that despite all the perplexities, decisions are regularly made in education based on an informal mixture of recently collected information, prior information, and the preferences of the decision maker. If these decisions are to be rational in the sense that they are derived from a logical program for decision making that provides for the input of prior information, new data, and preferences, we claim that decision theory is required. There is no magic in the formal structure of decision theory. The theory contributes only what mathematics does to any problem; an orderly, systematic, and precise framework for formulating a problem, plus the economy of mathematical reasoning in tracing the consequences of the formulation.

The main difficulty in implementing decision theory arises from the necessity to quantify basically subjective or personalistic quantities, and this difficulty is real. However, we believe that within the framework of some of the newer and very highly structured modes of instruction, it is possible to provide relatively simple yet conclusively meaningful methods of decision analysis. The decision machinery that we shall build will be appropriate for a wide band of decision problems. However, we will consistently illustrate the ideas with examples from individualized instructional procedures.

2. Fundamental Ideas and Notation

The main purpose of this section is to illustrate the rudiments of decision analysis with a variety of straightforward examples. Thus, we begin Section 2.1 with a rather trivial example. In later sections, this example will be modified and extended to make it both more realistic and interesting.

2.1 The Basic Example

The environment of our example is very simple. A student has either mastered or not mastered the topics in the current unit of his individually prescribed instructional program. The state of the student being a nonmaster will be denoted by θ_1 and the state of his being a master will be denoted by θ_2. Two actions are open to the decision maker, who could be the teacher or even, in some arrangements, the student himself. The decision may be to retain the student for additional work at the present level of his prescribed program. This "retain" decision will be denoted by d_1. The complement of this decision is to advance the student to the next unit. This decision will be denoted by d_2.

We now identify the three basic inputs into the decision analysis.

(a) Prior information. In Individually Prescribed Instruction, a student begins a unit of instruction only when he is deemed to be prepared for that unit. For this reason, the variation in posttest results tend to be relatively small, except to the extent that they are due to sampling variation. Therefore, the main input of prior information will involve our beliefs about the relative success of the instructional unit with qualified entrants. Thus, on the basis of an examination of the success which other students have had on this training unit, and before administering a test at the end of the present unit, the decision maker assigns prior probabilities to the two possible states.

Table 1

State of the Student

		θ_1 (nonmaster)	θ_2 (master)
Prior Probability	$\rho(\theta)$	.4	.6

This information is relevant to the decision under consideration and to ignore it is to waste useful facts.

The symbol $\rho(\theta)$ denotes the prior probability mass function. This symbol will continue to play the same role throughout our discussion.

(b) Experimental plan. The decision maker may ask the student one test question to learn more about which state prevails. The result of this short test will be denoted by X. If the student answers correctly, $X = 1$, and if he answers incorrectly, $X = 0$. The probability assignments to these two outcomes, if the mastery state is known, is given in the following short table where $p(x|\theta) = Pr[X = x|\theta]$ is the probability mass function of correct ($x = 1$) and incorrect ($x = 0$) responses given that the mastery level is known to be θ.

Table 2

$p(x|\theta)$

	$x = 0$ (wrong)	$x = 1$ (right)
θ_1 (nonmaster)	.8	.2
θ_2 (master)	.2	.8

This table states that if a person is a master, then the probability that he will give a correct answer is .8 and the probability that he will give a wrong response is .2. If he is a nonmaster, the probabilities are reversed. Note that those are probabilities of experimental outcomes <u>given</u> the true state of the student (master, nonmaster).

(c) Preferences. The decision maker can make one of the two types of errors. If he retains the student at the current level when, in fact, the student is a master, the student will probably repeat the current unit with only minimal gain. On the other hand, if the student is advanced when he has not mastered the topics on the current level, ultimately he may have to repeat both the current level and the one to which he had prematurely been advanced. With these facts in mind, the decision maker designates the nonnegative loss function $L(d, \theta)$ defined in Table 3.

Table 3

$L(d, \theta)$

	θ_1 (nonmaster)	θ_2 (master)
d_1 (retain)	0	1
d_2 (advance)	2	0

In specifying a loss function of this type, the decision maker assumes that no loss occurs if a correct decision is made and that the loss associated with advancing a nonmaster is twice that associated with retaining a master. We do not suggest that this simple loss function is appropriate in all or even any situations. However, the simple assumptions seem realistic enough to maintain our attention for a while.

Three reasonable decision rules for selecting decisions d_1 or d_2, after observing the test score, are possible.

Table 4

Rule	x = 0 (wrong)	x = 1 (right)
$\delta_1(x)$	d_1	d_1
$\delta_2(x)$	d_1	d_2
$\delta_3(x)$	d_2	d_2

In this trivial example, decision rules $\delta_1(x)$ and $\delta_3(x)$ really tell us to ignore prior information and the current test result; in the case of $\delta_1(x)$ to take action d_1 (retain) invariably, and in the case of $\delta_3(x)$ to stick with action d_2 (advance). Nevertheless, each is a serious candidate because certain loss structures and prior probability distributions could make one of them the preferred decision rule.

There are two ways that the analyses carried on within the black box labeled "decision process" have historically been organized. The first, which is called <u>normal form analysis</u>, involves a three step process. (1) Compute the average or expected loss for each ordered pair (δ_i, θ_j) composed of a decision rule δ_i and a state parameter value θ_j. This averaging is performed with respect to the probability distribution of possible experimental outcomes, $p(x|\theta_j)$.

The expected value that emerges from this computation is called a <u>risk function</u> and it is denoted by $R(\delta_i, \theta_j)$. As the notation suggests, it is a function of both the decision rule δ_i and the state parameter θ_j. In symbols, we have

$$R(\delta_i, \theta_j) = \sum_x L[\delta_i(x), \theta_j]p(x|\theta_j)$$

where i indexes the various decision rules and j identifies the various values of the state parameter θ. Once we have the risk function, we have the expected loss for each possible (δ_i, θ_j) pair, and can possibly decide that a decision rule is good (bad) if its risk is small (large) for all values of θ_j. (2) This kind of analysis, however, is typically inconclusive so we must compute the average or expected value of the risk function for each decision rule. This averaging is done with respect to the prior probability distribution of the state parameter $\rho(\theta)$. This

expected value is called the Bayes risk function and will be denoted by $r(\delta_i, \rho)$, where the symbol ρ is inserted to stress that the Bayes risk depends on the prior probability distribution. This recipe may be expressed in symbols as

$$r(\delta_i, \rho) = \sum_{\theta} R(\delta_i, \theta_j)\rho(\theta_j) \quad .$$

The effect of this averaging is to weight the effect of a decision rule highly for those values of θ_j that we think, apriori, are highly probable and, hence, important to consider. (3) With this computation completed for each decision rule, we select and subsequently use the decision rule that has the smallest Bayes risk.

Now let us carry out this three step normal form analysis using the losses and prior probabilities from our example. First, we will compute the risk of decision rule δ_2 when the true state of the person is θ_1. Recall that with δ_2, we retain the student if $x = 0$ and advance him if $x = 1$. The risk is

$$R(\delta_2, \theta_1) = \sum_{x} L[\delta_2(x), \theta_1]p(x|\theta_1)$$

$$= L[\delta_2(0), \theta_1]p(0|\theta_1) + L[\delta_2(1), \theta_1]p(1|\theta_1) \quad ,$$

i.e., the risk of using decision rule δ_2 when the true state is θ_1 is the simple average or expected loss for the (δ_2, θ_1) pair. The expectation is performed with respect to the probability distribution $p(x|\theta_1)$ of the two possible test scores $x = 0$ and $x = 1$, when it is given that the student is a nonmaster (i.e., $\theta = \theta_1$). Let us evaluate this risk. From Table 4, we know that $\delta_2(0) = d_1$ and $\delta_2(1) = d_2$. Thus,

$$R(\delta_2, \theta_1) = L(d_1, \theta_1)p(0|\theta_1) + L(d_2, \theta_1)p(1|\theta_1) \quad .$$

From Table 3, we have $L(d_1, \theta_1) = 0$ and $L(d_2, \theta_1) = 2$; from Table 2, we have $p(0|\theta_1) = .8$ and $p(1|\theta_1) = .2$. Thus, using the symbol * to indicate multiplication, we have

$$R(\delta_2, \theta_1) = 0 * .8 + 2 * .2 = .4 ,$$

i.e., when the true situation is θ_1, decision rule δ_2 has a risk (expected or average loss) of .4. Similar computations have been made for the remaining (δ_i, θ_j) pairs. The computations have been summarized in Table 5.

Table 5

$$R(\delta_i, \theta_j) = \sum_x L[\delta_i(x), \theta_j]p(x|\theta_j)$$

State	Result	$p(x\|\theta)$	$\delta_1(x)$	$\delta_2(x)$	$\delta_3(x)$
θ_1	x = 0	.8	0 * .8	0 * .8	2 * .8
	x = 1	.2	0 * .2	2 * .2	2 * .2
		$R(\delta_i, \theta_1)$	0	.4	2.0
θ_2	x = 0	.2	1 * .2	1 * .2	0 * .2
	x = 1	.8	1 * .8	0 * .8	0 * .8
		$R(\delta_i, \theta_2)$	1.0	.2	0

If we knew that θ_1 (nonmaster) prevailed, decision rule $\delta_1(x)$ which always retains the student would minimize the risk function. It has an expected loss of zero which is as good as one can do. If we knew that θ_2 (master) prevailed, decision rule $\delta_3(x)$ which always advances the student would minimize the risk function. Again this rule would have no risk. Of course, if we knew which state prevailed there would be no need to apply

decision analysis to the problem. At this point, none of the three decision rules are judged to be uniformly superior, i.e., superior for every state of nature. If one were, we would certainly adopt it, but none is, so we must find some way of choosing the best decision rule. We, therefore, move to Step (2) in normal form analysis before identifying the winner. The computation of the Bayes risk must now be made for each of the three decision rules. The idea is simply this. We do not know the true state of the person, but we do have a prior opinion concerning the true state. Therefore, it makes sense to average the risk for each decision rule with respect to our prior opinions, in effect, to put more weight on those values of θ that seem more probable to us. For example, to compute the Bayes risk $r(\delta, \rho)$ for δ_2 (the average risk with respect to our prior probabilities for θ_1 and θ_2), we compute the following:

$$r(\delta_2, \rho) = R(\delta_2, \theta_1)\rho(\theta_1) + R(\delta_2, \theta_2)\rho(\theta_2) \quad .$$

Substituting values from Table 5 and Table 1, we have

$$(.4 * .4) + (.2 * .6) = .28 \quad .$$

The computations of the Bayes risk for δ_1 and δ_3 are also easily made and are given, together with those for δ_2, in Table 6.

Table 6

θ	$\rho(\theta)$	$R(\delta_1, \theta)\rho(\theta)$	$R(\delta_2, \theta)\rho(\theta)$	$R(\delta_3, \theta)\rho(\theta)$
θ_1	.4	0 * .4	.4 * .4	2 * .4
θ_2	.6	1 * .6	.2 * .6	0 * .6
	$r(\delta_i, \rho)$	.6	.28	.80

The final step in normal form analysis directs us to select action rule δ_2 because it has the smallest Bayes risk. Thus, we have stated and exhibited in detail a precise and coherent procedure for decision-making in the presence of uncertainty. We have further demonstrated a simple application in the context of Individually Prescribed Instruction. To see that the choice of decision rule really depends on prior probabilities, the reader should redo step (3) in the analysis first with $\rho(\theta_1) = .1$, $\rho(\theta_2) = .9$, and then with $\rho(\theta_1) = .9$, $\rho(\theta_2) = .1$.

A second way of organizing the analysis within the decision process black box is called <u>extensive form analysis</u>. Since this type of analysis has some computational advantages over the normal form, our subsequent illustrations will, with one exception, employ extensive form analysis.

Extensive form analysis also involves a three step process. (1) Determine the posterior probability distribution of the state parameter θ, given the experimental result x. That is, we must determine $p(\theta|x) = p(x|\theta)\rho(\theta)/p(x)$, where $p(x) = \sum_{\theta} p(x|\theta)\rho(\theta)$ is the unconditional probability mass function of X. The posterior distribution $p(\theta|x)$ summarizes our knowledge and beliefs about θ, incorporating both our prior beliefs and the sample information. (2) Compute the expected value of the loss function for each decision rule with respect to the posterior distribution of θ. That is, we must compute $\sum_{\theta} L(\delta_1, \theta)p(\theta|x)$ for each decision rule. (3) Select the decision that will yield the smallest posterior expected loss in Step (2).

The advantage of extensive form analysis arises in Step (2) and is a bit hard to appreciate when expressed only in words. The heart of the matter is that Step (2) does <u>not</u> have to be carried out for every possible value of X. If we adopt this system, we can wait and perform Step (2) only

for the result x that is actually observed. Once x is observed, the decision rule δ_i specifies the decision d_i to be made, and the losses under each state θ can be taken immediately from Table 3.

The amazing thing is that under very general conditions, normal and extensive form analysis will lead to the same decision. This point will be illustrated in our example. Later, a mathematical argument will be presented for those who can only be persuaded by such demonstrations (Raiffa and Schlaifer, 1961, p. 15).

The first step in extensive form analysis requires us to determine the form of the posterior probability distribution of θ . For our example, we will do more and actually exhibit the two possible posterior distributions. But first, we need to exhibit the joint probability distribution of X and θ, and the marginal distribution of X . These probabilities are given in Table 7 where the entries in the body of the table are the joint probabilities, and the entries in the margins are the marginal probabilities. In computing Table 7, we have used $\rho(\theta)$ from Table 1 and $p(x|\theta)$ from Table 2.

Table 7

$p(x|\theta)\rho(\theta) = p(x, \theta)$

	$x = 0$	$x = 1$	$\rho(\theta)$
θ_1	.32	.08	.40
θ_2	.12	.48	.60
$p(x)$	.44	.56	

Then the posterior probability distribution of θ for given values of x is given by Bayes Theorem $p(\theta|x) = p(x, \theta)/p(x)$. These conditional probabilities for θ given x are summarized in Table 8. Note that the conditional distribution of θ given $x = 0$ is very different from that given $x = 1$.

Table 8

$p(\theta|x)$

	$x = 0$	$x = 1$
θ_1	.73	.14
θ_2	.27	.86
	1.00	1.00

The second step in extensive form analysis calls for the computation of the average or expected value of the loss function for each decision rule with respect to the posterior distribution of θ [i.e., we must compute $\Sigma_j L(\delta_i, \theta_j)p(\theta_j|x)$]. We will carry out this computation for each of the three decision rules specified in Table 4. Suppose $x = 0$. Then from Table 8, we see that $p(\theta_1|x = 0) = .73$ and $p(\theta_2|x = 0) = .27$. If we adopt δ_1, then we shall make decision d_1 when $x = 0$ (see Table 4). Thus, if $\theta = \theta_1$, our loss will be zero, and if $\theta = \theta_2$, our loss will be one (see Table 3). Therefore, our average or expected loss given $x = 0$ is

$$\mathcal{E}\{L[\delta_1(0), \theta]|x = 0\} = L[\delta_1(0), \theta_1]p(\theta_1|x = 0) + L[\delta_1(0), \theta_2]p(\theta_2|x = 0)$$
$$= (0 * .73) + (1 * .27) = .27$$

as given in the first column of Table 9.

Table 9

$\mathcal{E}\{L[\delta_1(x), \theta]|x\}$

	x = 0			x = 1		
	δ_1	δ_2	δ_3	δ_1	δ_2	δ_3
θ_1	0 * .73	0 * .73	2 * .73	0 * .14	2 * .14	2 * .14
θ_2	1 * .27	1 * .27	0 * .27	1 * .86	0 * .86	0 * .86
	.27	.27	1.46	.86	.28	.28

In this table, similar computations are made for each of the three decision rules for both x = 0 and x = 1. Observe that decision rule $\delta_2(x)$ produces the smallest expected loss for each value of x and may be judged as the optimum decision rule. However, if x = 0, rule δ_1 is equally as good, and if x = 1, rule δ_3 is equally as good. More importantly, however, this observation indicates that the second and third steps can be simplified still further. Note that if x = 0, $\delta_1(x)$ and $\delta_2(x)$ lead to the same decision (d_1) and, consequently, must necessarily have the same expected loss. What this emphasizes is that once we know x, we are really interested only in the best decision (d_1 or d_2), rather than the decision rule (δ_1, δ_2, δ_3) that will lead to the best decision. Consequently, we need only compute the expected or average value of the loss function for the available decisions. That decision (not decision rule) with the smallest posterior loss will then be selected. This is done in Table 10.

Table 10

$\sum_{\theta} L(d, \theta)p(\theta|x)$

	x = 0		x = 1	
	d_1	d_2	d_1	d_2
θ_1	0 * .73	2 * .73	0 * .14	2 * .14
θ_2	1 * .27	0 * .27	1 * .86	0 * .86
	.27	1.46	.86	.28

Our action rule is the same as indicated previously; if we observe $x = 0$, we will take action d_1 (retain), and if we observe $x = 1$, we will take action d_2 (advance). This is, in effect, the same as adopting decision rule δ_2. As pointed out previously, in extensive form analysis, only that half of Table 10 which corresponds to the actual result observed ($x = 0$ or $x = 1$) needs to be calculated.

So far, we have illustrated two forms of analysis which we claim will lead to identical decisions. These forms combine prior information, experimental results, and preferences using the machinery of probability theory to trace the consequences of the inputs. The example that helped to illustrate these ideas was kept at a trivial level so that the arithmetic would not obscure the essence of the process.

2.2 The Formal Structure of Decision Theory

Although mathematical symbols may seem forbidding, they are an indispensable tool in conveying ideas precisely. Therefore, both for completeness and for precision we will retrace the key ideas of decision analysis relying on symbols rather than numbers. Although integral signs, $\int$, will be used, we remind our readers that this is merely a continuous analogue of the summation operator, Σ.

We let $L[\delta(x), \theta)]$ be a nonnegative loss function defined on the set of pairs of the parameter values, δ and θ . The symbol $\delta(x)$ is used to emphasize the fact that the decision rules usually depend on observed values of the random variable X; the observation being obtained to elicit information about θ . In some contexts it may help to think of θ as the mean of some probability distribution, and $\delta(x)$ as an estimate of θ . Then $L[\delta(x), \theta]$ is the nonnegative numerical value of the loss associated with using $\delta(x)$ as an estimate of the mean, when, in fact, the mean is θ .

In the first stage of normal form analysis, we need the expected or average value of this loss function, where expectation is taken with respect to the conditional distribution of X given θ . This expected value was referred to as the risk function and was denoted $R(\delta, \theta)$. That is

$$R(\delta, \theta) = \int L[\delta(x), \theta]p(x|\theta)dx .$$

For this integral, the symbol $p(x|\theta)$ denotes the probability density of the random variable X, given the value of the state parameter θ . If X is a discrete random variable, $p(x|\theta)$ is to be interpreted as a probability mass function, and integration is then replaced by summation. As the symbols suggest, the risk function is similar to the loss function. The major difference rests on the fact that the risk function no longer depends upon the observed value of our experimental variable X; X has been averaged out. This fact is emphasized by the absence of an x in our notation for the risk function, $R(\delta, \theta)$.

If one adopts a Bayesian view of statistics, and one is compelled to embrace this view if he accepts any of several comprehensive axiom systems for decision making, it becomes necessary to quantify previous

or collateral information about the state parameter θ, which we summarize in $\rho(\theta)$. In the second stage of normal form analysis, the expected or average value of the risk function from stage one is then calculated with respect to this prior distribution, $\rho(\theta)$. We referred to this expected risk as the Bayes risk associated with δ, and denoted it, $r(\delta, \rho)$. And so we have,

$$r(\delta, \rho) = \int R(\delta, \theta)\rho(\theta)d\theta$$

$$= \int\{\int L[\delta(x), \theta]p(x|\theta)dx\}\rho(\theta)d\theta \quad .$$

Then, the third and final stage of the normal form analysis consisted of selecting the decision rule which minimizes the Bayes risk.

As was pointed out above, extensive form analysis follows a slightly different route, but under rather general conditions, leads one to the same decision. In extensive form analysis, we begin by evaluating the posterior distribution of θ

$$p(\theta|x) = \frac{\rho(\theta)p(x|\theta)}{p(x)} \quad ,$$

and then determine the expected loss with respect to this posterior distribution. In the continuous case, this expected or average loss can be represented by the integral,

$$\mathcal{E}_{\theta|x}\{L[\delta(x), \theta]\} = \int L[\delta(x), \theta]p(\theta|x)d\theta \quad .$$

Naturally, in the discrete case, we merely replace the integration operation with that of summation. The decision making criterion in extensive form analysis is then to choose that decision rule $\delta(x)$, which minimizes this expected loss. The reader should take special note at this point, that

although this expected loss depends upon the random variable X, it needs to be evaluated only for that value x which is actually obtained.

In order to make clear the relationship between normal form and extensive form analyses, let us compare the decision criteria in the two cases. Under normal form analysis we are to choose the decision which minimizes the Bayes risk,

$$r(\delta, \rho) = \int\{\int L[\delta(x), \theta]p(x|\theta)dx\}\rho(\theta)d\theta \quad .$$

Now, if these integrals are suitably well behaved, we may interchange the order of integration and so write the Bayes risk in the form,

$$r(\delta, \rho) = \int\{\int L[\delta(x), \theta]p(x|\theta)\rho(\theta)d\theta\}dx \quad .$$

Since by Bayes theorem $p(x|\theta)\rho(\theta) = p(\theta|x)p(x)$, we may rewrite this last equation in the form

$$r(\delta, \rho) = \int\{\int L[\delta(x), \theta]p(\theta|x)d\theta\}p(x)dx \quad . \qquad (2.1)$$

The observant reader will of course already have recognized the integral in brackets as the expected loss which must be minimized under the extensive form approach. This equation illustrates the crucial difference between the two approaches. Extensive form analysis chooses the decision rule which minimizes the expected loss for the <u>particular value of x observed</u>. In contrast, normal form analysis chooses the decision rule which minimizes the average of those expected losses <u>for all possible values of x</u>. Clearly, if one particular decision rule minimizes the expected loss criterion of

extensive form analysis for every x, then the average of those expected losses under that decision rule must also be a minimum. In this case, it is clear that the decision taken under extensive form analysis will coincide precisely with that taken under normal form analysis. However, in those instances where extensive form analysis apparently leads to different decision rules depending upon the value of x observed, the equivalence of the two approaches may not be obvious.

Although the possible non-equivalence of these two approaches may seem to pose a dilemma for users of decision theory, in fact, it is a non-problem which has been set merely for the pedagogical purpose of underscoring some fundamental differences in the two approaches. Consider, if you will, the following. Since we are admitting for consideration all reasonable decision rules, we must allow that rule $\delta^*(x)$ which, for each x, minimizes the expected loss criterion of extensive form analysis. Since $\delta^*(x)$ will be selected by the extensive form approach irrespective of the value of x obtained, it will also be chosen under the normal form procedure. How do we construct $\delta^*(x)$? We do this in a straightforward operational manner: We use extensive form analysis for the x obtained and choose that <u>decision</u> (not decision rule) which minimizes the expected loss.

For a concrete example of the relationship between normal form and extensive form analysis, consider Table 11 which summarizes the expected losses under extensive form analysis for the three decision rules. This table is merely a modified version of Table 9.

Table 11

$\mathcal{E}_{\theta|x}\{L[\delta_i(x), \theta]\}$

	δ_1	δ_2	δ_3
x = 0	.27	.27	1.46
x = 1	.86	.28	.28

We found in Table 7 that the marginal distribution of X is given by pr(X = 0) = .44 and pr(X = 1) = .56. Thus, as indicated in Equation (2.1), the Bayes risk associated with decision δ_1 is given by the weighted average of the entries in the first column of Table 11, where the weights are the marginal probabilities of X . And so we have,

$$r(\delta_1, \rho) = (.27 * .44) + (.86 * .56) = .6$$

as we saw in Table 6.

2.3 Extensive Form Analysis with a Continuous Posterior

Let us now turn to an example of extensive form analysis which uses a continuous model density. This example is only a slight modification of that used previously; the primary difference being that we now assume that both the state parameter θ and the random variable X are continuous. It should be noted that this example is merely a reformulation of one considered by Hambleton and Novick (1972). As before, two decisions are open to the decision maker who is guiding a student through an ordered sequence of instructional units. At the end of each unit, the decision maker, based on his knowledge of the student's past performance, the performance of similar students, and current test results, must decide

to advance the student to the next unit in the sequence or to retain him at the present level. If the decision maker knows the student's mastery level (θ), he would be willing to advance the student if $\theta \geq \theta_o$ and to retain him if $\theta < \theta_o$. Thus, the number θ_o is a cutoff (or selection) point on the mastery scale with respect to the actions advance or retain.

In selecting θ_o, careful consideration of the objectives of the training program, and previous experience with the training and evaluation materials must prevail. If for example, θ_o is intended merely to give at least an even chance of completing the next lesson, θ_o might be set equal to that level of functioning which has historically had a 50% success rate on the next unit. If, on the other hand, the decision maker is very concerned about the ill-effects of the frustration of a poorly prepared student reading advanced material, perhaps θ_o ought to be somewhat higher. In any case, once θ_o is specified for the test, prior and collateral information about the student will be combined with the test result (x) for the purpose of estimating θ.

Assume that for this two-action (advance or retain) problem, the decision maker specifies a threshold (or step) loss function which can be described by the following table: (Compare with Table 3)

Table 12

$L(d_i, \theta)$

	$\theta < \theta_o$	$\theta \geq \theta_o$
d_1 (retain)	0	b
d_2 (advance)	c	0

where b and c are both nonnegative. In the literature, losses associated with falling into the lower left cell of this table are frequently referred to as arising from "false-positives", since the decision maker has wrongly presumed that the parameter θ lies in the region which has positive ethical value (i.e., has wrongly presumed that $\theta \geq \theta_o$). Similarly, losses associated with falling in the upper right cell are commonly referred to as arising from "false-negatives". The nonnegative numbers, b and c, reflect the cost of making these two types of errors. Because the decision envisaged in this example is rather local, affecting only one step in a program which is only a small part of the student's total learning experience, a massive effort to determine b and c exactly would seem inappropriate. In some cases, it might seem reasonable to assume that c/b = 2, if for example, a false positive results in repeating two steps in the sequence, as compared with the repetition of only one for a false negative.

Following the general scheme for extensive form analysis outlined earlier, our goal is to determine the action which will minimize the expected or average posterior loss

$$\int L(d_i, \theta)p(\theta|x)d\theta \quad .$$

This integral is equal to $b\{Pr[\theta \geq \theta_o|x]\}$ if i = 1 and is equal to $c\{Pr[\theta < \theta_o|x]\}$ if i = 2. Therefore, we may minimize this expected loss by making decision d_1 for those values of x such that

$$b\{Pr[\theta \geq \theta_o|x]\} < c\{Pr[\theta < \theta_o|x]\}$$

and by making decision d_2 for those values of x where

$$b\{Pr[\theta \geq \theta_o|x]\} > c\{Pr[\theta < \theta_o|x]\} \quad .$$

Although such a situation is unlikely in practice, when the two possible values of the integral are equal, we will be indifferent about which decision to take.

In many applications of this type, the range of test scores may be cut into a decision d_1 region and a decision d_2 region by considering the posterior distribution of θ[i.e., $p(\theta|x)$] as a function of x . Doing this, we see that the critical point dividing the two regions can be represented by that point x_o in the set of possible test results such that

$$c\{Pr[\theta < \theta_o|x_o]\} = b\{Pr[\theta \geq \theta_o|x_o]\}$$

or

$$c\{Pr[\theta < \theta_o|x_o]\} = b\{1 - Pr[\theta < \theta_o|x_o]\}$$

so that

$$\frac{b}{c + b} = Pr[\theta < \theta_o|x_o] \text{ .}^2 \qquad (2.2)$$

What this equation says is that if we consider the class of possible posterior distributions $\{p(\theta|x)\}$ to be indexed by x, and if we can find that member of this class which is identified by $x = x_o$, say, such that $Pr[\theta < \theta_o|x_o] = b/(c + b)$, then for all $x < x_o$, we will choose decision d_1 and for all $x > x_o$, we will select decision d_2 .

We will illustrate the computation of the cutting score x_o on our observation scale with an example. Suppose that the test score X has a normal distribution with unknown mean θ and <u>known</u> variance σ^2 . Further, suppose that all existing information about the parameter θ, which measures the mastery level of a certain skill, may be summarized by a normal distribution with mean τ and variance ϕ . Then, a simple application of Bayes theorem

---

[2] The uniqueness of the point x_o satisfying this equation is presumed.

yields a posterior distribution of θ which is normal and has a mean of $(\tau\sigma^2 + x\phi)/(\sigma^2 + \phi)$ and a variance of $\sigma^2\phi/(\phi + \sigma^2)$. Transforming this posterior distribution on θ into standard form, we see that

$$\Pr[\theta < \theta_o | x_o] = \Pr\left[z < \frac{\theta_o - \dfrac{\tau\sigma^2 + x_o\phi}{\sigma^2 + \phi}}{[\phi\sigma^2/(\sigma^2 + \phi)]^{\frac{1}{2}}} \,\middle|\, x_o \right] ,$$

where z has a normal distribution with mean zero and variance one. Therefore, the cutting score x_o on the observation scale may be determined by finding that point z_o in the standard normal distribution which has percentile rank $100[b/(c + b)]$, and then solving the equation

$$z_o = \frac{\theta_o - \dfrac{\tau\sigma^2 + x_o\phi}{\sigma^2 + \phi}}{[\phi\sigma^2/(\sigma^2 + \phi)]^{\frac{1}{2}}}$$

for x_o . Thus,

$$x_o = \frac{\phi + \sigma^2}{\phi} \{\theta_o - z_o[\phi\sigma^2/(\phi + \sigma^2)]^{\frac{1}{2}}\} - \frac{\tau\sigma^2}{\phi} . \qquad (2.3)$$

In order to convey some feeling for how this loss structure and the normal data and normal prior distributions interact to produce cutting scores, Table 13 has been provided. In this table, the desired proficiency level $\theta_o = 75$ and the prior mean $\tau = 80$.

As we would expect with ϕ and σ^2 held constant, the cutting score x_o increases with c/b, the relative loss for a false positive as compared with a false negative error. From an intuitive point of view, this relationship makes sense. As c/b increases, false positive errors become relatively more expensive than those of the false negative variety. Consequently, our

Table 13

Loss Constants		Variances		Posterior	Posterior Cutting	Critical Posterior
c	b	Prior ϕ	Data σ^2	Variance	Score x_o	Mean $\frac{80\,\sigma^2 + x_o\phi}{\sigma^2 + \phi}$
1	9	9.0	9.0	4.5	64.57	72.28
1	9	9.0	16.0	5.8	57.58	71.93
1	9	9.0	25.0	6.6	48.67	71.71
1	9	16.0	9.0	5.8	67.39	71.93
1	9	16.0	16.0	8.0	62.76	71.38
1	9	16.0	25.0	9.8	56.94	71.00
1	9	25.0	9.0	6.6	68.72	71.71
1	9	25.0	16.0	9.8	65.24	71.00
1	9	25.0	25.0	12.5	60.95	70.47
3	7	9.0	9.0	4.5	67.77	73.89
3	7	9.0	16.0	5.8	62.61	73.74
3	7	9.0	25.0	6.6	56.01	73.65
3	7	16.0	9.0	5.8	70.22	73.74
3	7	16.0	16.0	8.0	67.03	73.52
3	7	16.0	25.0	9.8	62.99	73.36
3	7	25.0	9.0	6.6	71.36	73.65
3	7	25.0	16.0	9.8	69.11	73.36
3	7	25.0	25.0	12.5	66.29	73.14
5	5	9.0	9.0	4.5	70.00	75.00
5	5	9.0	16.0	5.8	66.11	75.00
5	5	9.0	25.0	6.6	61.11	75.00
5	5	16.0	9.0	5.8	72.19	75.00
5	5	16.0	16.0	8.0	70.00	75.00
5	5	16.0	25.0	9.8	67.19	75.00
5	5	25.0	9.0	6.6	73.20	75.00
5	5	25.0	16.0	9.8	71.80	75.00
5	5	25.0	25.0	12.5	70.00	75.00
7	3	9.0	9.0	4.5	72.23	76.11
7	3	9.0	16.0	5.8	69.61	76.26
7	3	9.0	25.0	6.6	66.21	76.35
7	3	16.0	9.0	5.8	74.16	76.26
7	3	16.0	16.0	8.0	72.97	76.48
7	3	16.0	25.0	9.8	71.39	76.64
7	3	25.0	9.0	6.6	75.04	76.35
7	3	25.0	16.0	9.8	74.49	76.64
7	3	25.0	25.0	12.5	73.71	76.86
9	1	9.0	9.0	4.5	75.43	77.72
9	1	9.0	16.0	5.8	74.64	78.07
9	1	9.0	25.0	6.6	73.55	78.29
9	1	16.0	9.0	5.8	76.99	78.07
9	1	16.0	16.0	8.0	77.24	78.62
9	1	16.0	25.0	9.8	77.43	79.00
9	1	25.0	9.0	6.6	77.68	78.29
9	1	25.0	16.0	9.8	78.36	79.00
9	1	25.0	25.0	12.5	79.05	79.53

decision maker decreases his chances of making false positive errors by increasing the cutting score, x_o .

Since the mean of the posterior distribution of θ is a linear function of x, we may reformulate the question of the critical cutting score x_o into a question concerning the critical posterior mean $\mu_o = (\tau\sigma^2 + x_o\phi)/[\phi\sigma^2/(\sigma^2 + \phi)]^{\frac{1}{2}}$ where x_o is determined by Equation (2.3). In this case, we solve the equation

$$z_o = \frac{\theta_o - \mu_o}{[\phi\sigma^2/(\sigma^2 + \phi)]^{\frac{1}{2}}}$$

for μ_o . And analogous to Equation (2.3) above for x_o, we have

$$\mu_o = \theta_o - z_o\,[\phi\sigma^2/(\sigma^2 + \phi)]^{\frac{1}{2}} .$$

In preparing Table 13, we also calculated values of μ_o which appear in the last column. These results can be understood by referring to Figure 2.1. Our procedure says that the critical posterior distribution of θ must be such that $100[b/(c + b)]$ percent of the probability lies below θ_o . In order to maintain this constant percentage, as the posterior variance increases, the critical posterior mean of θ must necessarily decrease for $\frac{b}{(c + b)} > .5$ and increase for $\frac{b}{(c + b)} < .5$. That is, as the posterior variance of θ increases, the critical posterior mean moves away from θ_o . From an intuitive point of view, this makes sense. It implies that as a decision maker becomes increasingly uncertain about the posterior mean as an estimate of θ, he becomes more cautious, moving his critical mean in the direction of the less costly errors.

The effects of changes in parameter values on x_o is a bit more complicated since we must consider not only the posterior variance, but also the ratio of the prior variance (ϕ) to the variance of sampling

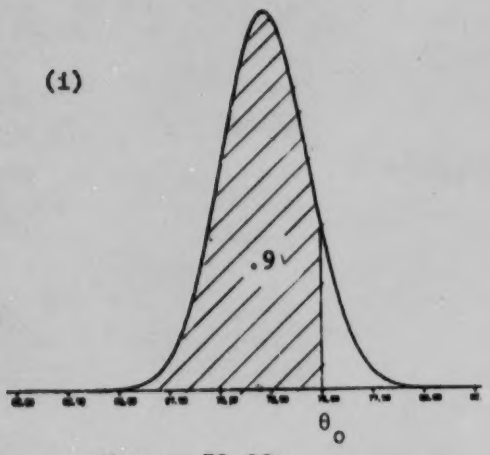

Mean = 72.28
Variance = 4.5

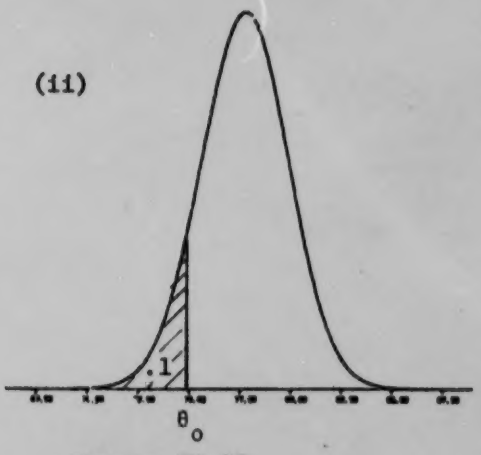

Mean = 77.72
Variance = 4.5

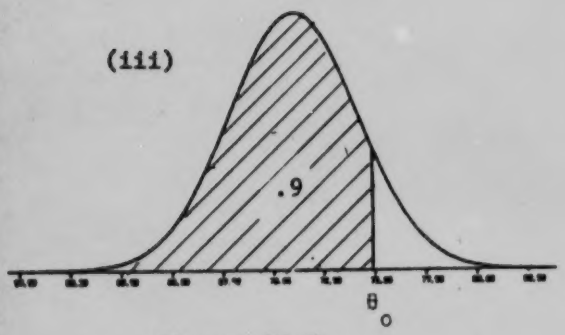

Mean = 71.0
Variance = 9.8

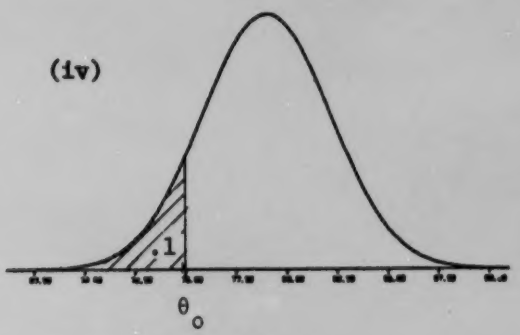

Mean = 79.00
Variance = 9.8

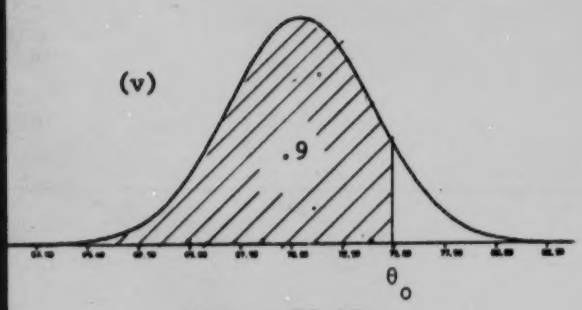

Mean = 70.47
Variance = 12.5

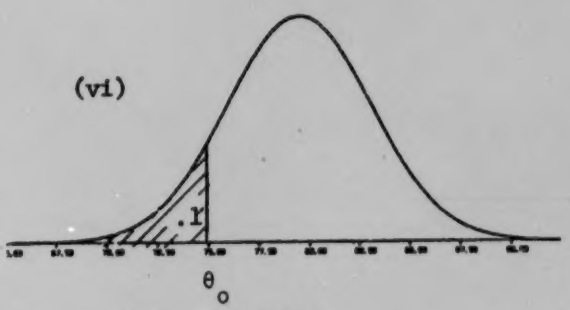

Mean = 79.53
Variance = 12.5

Figure 2.1. This figure illustrates the necessary change in the critical posterior mean as the posterior variance increases. In the figure, we let $b/(b + c) = .9$ in (i), (iii), and (v). We let $b/(b + c) = .1$ in (ii), (iv), and (vi). The figure illustrates that for $b/(b + c) > .5$, the critical posterior mean of θ must decrease with increasing variance, while for $b/(b + c) < .5$ it must increase.

distribution of x (i.e., σ^2). One can get some notion of what is going on by examining the definition of μ_o .

$$\mu_o = \frac{\tau\sigma^2 + x_o\phi}{\sigma^2 + \phi}$$

Since μ_o is the weighted average of τ and x_o, for given μ_o, x_o must increase as the ratio ϕ/σ^2 increases. This relationship is clearly indicated in Table 13.

At this point let us look at a numerical example of some of the theory we have developed in the last few pages. Assume that we have the following situation:

(1) $\rho(\theta) \sim N(80, 25)$ (Prior Distribution)

(2) $p(x|\theta) \sim N(\theta, 16)$ (Distribution of Test Scores)

(3)

$L(d_i, \theta)$ Loss Function

	$\theta < \theta_o$	$\theta \geq \theta_o$
d_1 (retain)	0	1
d_2 (advance)	2	0

(4) $\theta_o = 75$ (Pass Level)

Before collecting our observation on this subject, we can write the posterior probability density function $p(\theta|x)$ as a function of x . Thus,

$$p(\theta|x) \sim N\left[\frac{80(16) + 25x}{16 + 25} \;,\; \frac{25(16)}{25 + 16}\right] .$$

Since this is a completely general description of the posterior distribution no matter what value x is observed, it should be possible to determine those values of x which lead us to decision d_1 and those which lead us to d_2. This is what we did symbolically before when we determined the critical score x_o. We defined x_o as that point such that if x happens to fall below x_o, then applying our criterion of minimum expected loss, we would be led to choose d_1. And of course, if x falls above x_o, we would choose d_2. The previous theory tells us that the next step is to calculate the expected loss under each decision. Doing this we find:

$$\mathcal{E}_{\theta|x}[L(d_1, \theta)] = \int_{75}^{\infty} 1 \; p(\theta|x)d\theta$$

and

$$\mathcal{E}_{\theta|x}[L(d_2, \theta)] = \int_{-\infty}^{75} 2 \; p(\theta|x)d\theta \; .$$

By equating these two expected values symbolically and solving for that value x for which equality obtains, we previously found that

$$x_o = \frac{\phi + \sigma^2}{\phi} \{\theta_o - z_o [\phi\sigma^2/(\phi + \sigma^2)]^{\frac{1}{2}}\} - \tau\frac{\sigma^2}{\phi}$$

where z_o is the $\left(\frac{100b}{c + b}\right)$th percentile of the unit-normal distribution. Substituting into the equation for x_o, we find:

$$x_o = \frac{25 + 16}{25} \{75 - (-.43) [25(16)/(25 + 16)]^{\frac{1}{2}}\} - 80\,\frac{16}{25}$$

$$= 74.003 \; .$$

Therefore, we are certain that if we observe an x (test score) which is greater than 74.003, extensive form analysis will lead us to advance the student. Similarly, if the observed test score is less than 74.003, extensive form analysis will cause us to retain the student for additional training. Although, obtaining a score of precisely 74.003 will be somewhat ambiguous

since the two decisions will have the same expected loss, in practice this will not be a problem. We will seldom find a test score which falls precisely on an indifference point. Suppose that we observe x = 73. Since 73 is less than $x_o = 74.003$, we can be certain that extensive form analysis will lead us to retain the student (decision d_1). However, for the sake of those who doubt mathematical arguments, we will calculate the expected losses for comparison.

$$\mathcal{E}[L(d_1, \theta)|x = 73] = \int_{75}^{\infty} 1 \; p(\theta|x = 73)d\theta = 1 - Pr(\theta < 75|x = 73)$$

$$\mathcal{E}[L(d_2, \theta)|x = 73] = \int_{\infty}^{75} 2 \; p(\theta|x = 73)d\theta = 2[Pr(\theta < 75|x = 73)] \; .$$

Since

$$p(\theta|x = 73) \sim N(75.73, 9.756),$$

we have

$$Pr(\theta < 75|x = 73) = .409 \; ,$$

and so

$$\mathcal{E}[L(d_1, \theta)|x = 73] = .591$$

and

$$\mathcal{E}[L(d_2, \theta)|x = 73] = .818 \; .$$

Clearly, we would choose d_1 and retain the student as predicted.

A more realistic model than that just described would recognize the fact that in educational settings the model variance as well as the mean is usually unknown. In most situations, the decision maker will have some information concerning the variation and the region in which the observations will fall; total ignorance would preclude even the proper

choice of a measuring instrument. However, except in very special situations, one's knowledge about the model variance σ^2 is typically not sufficiently precise to warrant the application of the known variance model.

The solution to the inference part of the unknown mean and unknown variance problem has been provided by Novick and Jackson (1974). Beginning with a normal model density with location parameter θ and variance parameter ϕ, and a gamma-normal prior density proportional to

$$b_o(\theta, \phi) \propto \frac{1}{\phi^{\frac{m}{2}+1}} \exp\left\{\frac{m(\theta - w.)^2 + R^2}{-2\phi}\right\},$$

they are able to show that

$$t(\theta) = \frac{[m(m+1)]^{\frac{1}{2}}[\theta - (mw. + x)/(m+1)]}{[R^2 + m(x - w.)^2/(m+1)]^{\frac{1}{2}}} \qquad (2.4)$$

has a student's t distribution on m degrees of freedom. In this equation, $R^2/(m+1)$ is the prior modal estimate of ϕ; w. is the prior modal estimate of θ given that $\phi = R^2/(m+1)$; and, the parameter m is a weight factor which describes the decision maker's degree of confidence in his estimates. For the details of this development, the reader is referred to Novick and Jackson (1974, Chapter 7).

From Equation (2.4), we see that $t(\theta)$ is linear in θ. Thus, $Pr(\theta < \theta_o|x) = Pr[t(\theta) < t(\theta_o)|x]$. And by using Equation (2.2), we can partition the observation scale into two disjoint regions; one which will lead to decision d_1 and the other which will lead to decision d_2. The process is practically the same as in the known variance case. The only important change is that we now use the t table with m degrees of

freedom instead of a table of the unit normal distribution. Thus, we determine the point t_o which has percentile rank $100b/(c + b)$ and then solve the equation

$$t_o = \frac{[m(m + 1)]^{\frac{1}{2}}[\theta_o - (mw. + x_o)/(m + 1)]}{[R^2 + m(x_o - w.)^2/(m + 1)]^{\frac{1}{2}}} \qquad (2.5)$$

for x_o . Although the process of solving this equation will ordinarily lead to two results, only one of the results will solve Equation (2.5) as stated. The other will be associated with Equation (2.5) with $-t_o$ on the left-hand side.

2.4 An Example with a Binomial Model

Let us examine this same basic problem using a different, and in some ways, a more general model. For pedagogical purposes, this analysis will initially be of the normal form variety, for we will actually exhibit the Bayes risk function. Later we will redo the analysis using the much simpler extensive form. Rather than assuming that the mastery level θ takes on values in an interval, let us return to the original situation where only two values are possible. As before, we use the symbol θ_1 to denote the class of nonmasters and θ_2 the class of masters at any point in time. On the basis of prior information about the student and his training, the decision maker formulates a prior probability distribution on the two-point state space, $\rho(\theta_1) = 1 - p$ and $\rho(\theta_2) = p$, where $0 \leq p \leq 1$. Thus, p represents the "prior" probability that a given student is a master. Clearly, if p were equal to zero or one, no uncertainty would exist and no decision problem would remain. As before, two actions are open to the decision maker: d_1, declare the student a nonmaster and retain him at his present level or, d_2, declare the student a master and advance him to the next level. Analogous to Table 3, our decision maker adopts the familiar threshold loss function:

Table 14

$L(d_i, \theta_j)$

	θ_1	θ_2
d_1 (retain)	0	b
d_2 (advance)	c	0

where b and c are nonnegative.

It will be convenient in this and many applications to conceptualize a hypothetical population of tasks for which the mastery judgments are relevant. This done, we define α and β as the conditional probabilities of acceptably completing a randomly selected task from this population, given that the student is a master or nonmaster, respectively. In most applications α will be large and β will be small. Our decision maker plans to construct a mastery test by selecting t tasks from our hypothetical population. For the purposes of this example, we will assume that the tasks are experimentally independent and that the probability of success on each task depends only upon the mastery class to which the student belongs. We denote by X the discrete random variable associated with the number of tasks successfully completed. These assumptions imply that, given θ, the random variable X has a binomial mass function given by

$$p(x|\theta) = \begin{cases} \binom{t}{x}\beta^x(1-\beta)^{t-x} & \text{for } \theta = \theta_1 \\ \binom{t}{x}\alpha^x(1-\alpha)^{t-x} & \text{for } \theta = \theta_2 \end{cases} \tag{2.6}$$

where $x = 0, 1, 2, \ldots, t$.

If we are to follow the plan of normal form analysis previously outlined, we should now write down all reasonable decision rules. With $t + 1$ possible outcomes (i.e., $x = 0, 1, 2, \ldots, t$), an exhaustive list of decision rules for even moderate t, would be cumbersome, indeed. However, in the opinion of these writers the totality of reasonable decision rules for this problem can be summarized by the relationship

$$\delta_s(x) = \begin{cases} \text{retain} & \text{if } x < s \\ \text{advance} & \text{if } x \geq s \end{cases} \tag{2.7}$$

for $s = 0, 1, 2, \ldots, t + 1$. What this set of decision rules boils down to is the following: Choosing a decision rule $\delta_s(x)$ is equivalent to choosing a cutting score s on the test score scale such that, if the observed number correct (x) lies below s, the student is retained at the present level, otherwise he is advanced. Selection of the decision rule $\delta_{10}(x)$, for example, would lead one to retain the student if he completes 9 items or less correctly, and advance him, if he obtains a total score of at least 10.

For this problem, the risk $R(\delta_s, \theta_j)$ associated with each (δ_s, θ_j) pair can now be symbolically represented by

$$R(\delta_s, \theta_j) = \mathcal{E}_{x|\theta}\{L[\delta_s(x), \theta_j]\} \quad .$$

Inserting the loss function and model density into this relationship, we find that

$$R(\delta_s, \theta_1) = \sum_{x=s}^{t} c\binom{t}{x}\beta^x(1 - \beta)^{t - x} \tag{2.8a}$$

and

$$R(\delta_s, \theta_2) = \sum_{x=0}^{s-1} b\binom{t}{x}\alpha^x(1 - \alpha)^{t - x}$$

$$= b\left[1 - \sum_{x=s}^{t} \binom{t}{x}\alpha^x(1 - \alpha)^{t - x}\right] . \tag{2.8b}$$

Thus, the Bayes risk for each δ_s may be represented by

$$r(\delta_s, \rho) = \mathcal{E}_\theta [R(\delta_s, \theta)]$$

$$= (1 - p)c \sum_{x=s}^{t} \binom{t}{x}\beta^x(1 - \beta)^{t - x} + bp[1 - \sum_{x=s}^{t} \binom{t}{x}\alpha^x(1 - \alpha)^{t - x}]$$

$$= bp + \sum_{x=s}^{t} [(1 - p)c\binom{t}{x}\beta^x(1 - \beta)^{t - x} - bp\binom{t}{x}\alpha^x(1 - \alpha)^{t - x}] . \tag{2.9}$$

Now it is clear that if the expression inside the summation sign has a structure such that it is positive for all x less than some value x_o and negative for all x greater than x_o, we could minimize the Bayes risk, $r(\delta_s, \rho)$, by choosing $\delta_s(x)$ where s is the smallest integer such that $s \geq x_o$. In fact, this expression has the necessary form. To see this, all one needs to do is set the expression of interest equal to zero and solve

$$c(1 - p)\binom{t}{x_o}\beta^{x_o}(1 - \beta)^{t - x_o} - bp\binom{t}{x_o}\alpha^{x_o}(1 - \alpha)^{t - x_o} = 0$$

for the value(s) of x_o . A few routine manipulations yield the root

$$x_o = \frac{\ell n(c/b) + \ell n[(1 - p)/p] + t\ \ell n[(1 - \beta)/(1 - \alpha)]}{\ell n[(1 - \beta)/(1 - \alpha)] + \ell n[\alpha/(\beta)]} . \tag{2.10}$$

Since there exists only one root x_o, the expression, thought of as a continuous function of x, can cross the x axis at only one point. Therefore, it must be true that there is exactly one region of the x scale where the expression is positive and exactly one region where it is negative. Although this argument assures us that there is exactly one root, it does not reveal in which region the expression is positive and in which it is negative. We, therefore, re-examine the expression in brackets in Equation (2.9).

For notational convenience, we define Q to be the expression of interest. Thus,

$$Q = (1 - p)c\binom{t}{x}\beta^{x}(1 - \beta)^{t - x} - bp\binom{t}{x}\alpha^{x}(1 - \alpha)^{t - x} .$$

Also for convenience, we define a quantity R by the equation

$$R = \frac{(1 - p)c\binom{t}{x}\beta^{x}(1 - \beta)^{t - x}}{bp\binom{t}{x}\alpha^{x}(1 - \alpha)^{t - x}}$$

$$= \left(\frac{c}{b}\right) \left(\frac{1 - p}{p}\right) \left(\frac{1 - \beta}{1 - \alpha}\right)^{t} \left[\frac{\beta(1 - \alpha)}{\alpha(1 - \beta)}\right]^{x} .$$

Our purpose in the following development is to determine that region in which Q is negative. As a first step, we assume that $Q < 0$ and determine the implications of that assumption. If $Q < 0$, then by merely manipulating the inequality, $Q < 0$, we see that $R < 1$ and, therefore, $\ell n\ R < 0$. Thus,

$$\ell n\ R = \ell n \left(\frac{c}{b}\right) + \ell n \left(\frac{1 - p}{p}\right) + t\ \ell n \left(\frac{1 - \beta}{1 - \alpha}\right) - x\ \ell n \left[\frac{\alpha(1 - \beta)}{\beta(1 - \alpha)}\right] < 0$$

and so,

$$x\ \ell n \left[\frac{\alpha(1 - \beta)}{\beta(1 - \alpha)}\right] > \ell n \left(\frac{c}{b}\right) + \ell n \left(\frac{1 - p}{p}\right) + t\ \ell n \left(\frac{1 - \beta}{1 - \alpha}\right) . \quad (2.11)$$

Comparing Equation (2.10) and Equation (2.11), we see that if $\ell n\{[\alpha(1 - \beta)]/[\beta(1 - \alpha)]\} > 0$, the condition that $Q < 0$ is satisfied whenever $x > x_o$. And if $\ell n\{[\alpha(1 - \beta)]/[\beta(1 - \alpha)]\} < 0$, the condition that $Q < 0$ is satisfied whenever $x < x_o$. But the condition that $\ell n\{[\alpha(1 - \beta)]/[\beta(1 - \alpha)]\} < 0$ is satisfied if and only if $[\alpha(1 - \beta)]/[\beta(1 - \alpha)] < 1$, which is equivalent to $\alpha < \beta$. Admitting that the foregoing is a bit confusing, we summarize the results in the following table.

	$x < x_o$	$x > x_o$
$\alpha < \beta$	$Q < 0$	$Q > 0$
$\alpha > \beta$	$Q > 0$	$Q < 0$

The implications of the statements in this table for the decision-making process are two.

(1) If $\alpha > \beta$, then the decision maker will minimize the Bayes risk $r(\delta_s, \rho)$ by choosing $\delta_s(x)$ where s is the smallest integer such that $s \geq x_o$.

(2) If $\alpha < \beta$, the decision function (2.7) is inappropriate. The condition $\alpha < \beta$ implies that the decision maker's model asserts that a nonmaster is more likely to get a particular item correct than a master. Such a model would certainly lead the decision maker to consider a decision function of the form

$$\delta_s(x) = \begin{cases} \text{advance} & \text{if } x < s \\ \text{retain} & \text{if } x \geq s \end{cases}$$

as more appropriate than Equation (2.7).

Since under normal conditions α will be considerably greater than β, implication number one above will usually apply.

Perhaps the most important characteristic of the expression for the critical or cutting test score, Equation (2.10), is its dependence on the probabilities α and β . When α and $1 - \beta$ are both near one, the cutting score will tend to be small. Crudely speaking, in this case it does not take many satisfactory performances to decide whether a student is a master or not. On the other hand as α and β approach one another, it becomes ever more costly to separate the masters and nonmasters. If s

is greater than t, the number of tasks in the test is too small to permit any judgment other than d_1 (retain). If $s \leq 0$, the only feasible decision is to advance the student. This might occur if p, the prior probability of state θ_2, is very close to one. Table 15 gives some examples of critical or cutting scores for some selected parameter values, with the loss function constants c = 2 and b = 1 in Table 14.

Returning now to familiar ground, we review the implications of the latest wrinkle in our decision making scheme in the context of our initial numerical example. The only structural difference between the present situation and that of Section 2.1 is that we now have t tasks instead of one. For the purposes of this example, let us assume that we have a test of length eight (i.e., t = 8). Then, following the procedure outlined for normal form analysis problems, we identify the values of the inputs to the black box.

(a) Specification of Prior Information. From Table 1, we see that our prior beliefs about θ may be summarized by $\rho(\theta_1) = .4$ and $\rho(\theta_2) = .6$. Thus, in the notation of this example, we have p = .6, or the odds we would be just willing to give that this student is a master without resort to current test score information are 3/2.

(b) Indicating the Experimental Plan. As we pointed out in the previous development of this example, the plan is for the decision maker to give the student a test composed of 8 tasks. On the basis of this test, the decision maker is to give the student a score which is equal to the number of tasks correctly answered. The assumptions are, of course, that the tasks are equally difficult and experimentally independent, given θ. The distribution of X given θ can therefore be described by the binomial mass function

Table 15

Cutting Scores

(Classify as a master (θ_2) if the number correct equals or exceeds the cutting score)

Prior Prob of master	Number of test tasks	Prob. of success for master (α), Prob of success for nonmaster (β)							
		α	β	α	β	α	β	α	β
p	t	.8	.2	.8	.1	.9	.2	.9	.1
.9	10	5		4		6		5	
	20	10		8		12		10	
	30	15		13		17		15	
.7	10	5		5		6		5	
	20	10		9		12		10	
	30	15		13		18		15	
.5	10	6		5		6		6	
	20	11		9		12		11	
	30	16		13		18		16	

$$p(x|\theta) = \begin{cases} \binom{t}{x}\beta^x(1-\beta)^{t-x} & \text{for } \theta = \theta_1 \\ \binom{t}{x}\alpha^x(1-\alpha)^{t-x} & \text{for } \theta = \theta_2 . \end{cases}$$

From Table 2, we see that for this example $\alpha = .8$ and $\beta = .2$. This states that if our student is a master, the probability that he will give an acceptable response to any task is $\alpha = .8$. If he is a nonmaster, this probability is $\beta = .2$.

(c) Specifying Preferences and Decision Rules. From Table 3, we see that in the notation of the previous theoretical development, our threshold loss may be described by $c = 2$ and $b = 1$. In words, this implies that we would be twice as unhappy (in terms of some measure of loss) if we were to advance a nonmaster than we would be if we retained a master--$L(d_2, \theta_1)/L(d_1, \theta_2) = 2$.

Following the procedure indicated in our general development, we identify 10 reasonable decision functions.

$$\delta_s(x) = \begin{cases} d_1 \text{ (retain)} & x < s \\ d_2 \text{ (advance)} & x \geq s \end{cases}$$

for $s = 0, 1, 2, \ldots, 9$.

The ramifications of these decisions are summarized in the following table.

Decision	Retain	Advance
$\delta_o(x)$	Never	Always
$\delta_1(x)$	if $x = 0$	if $x \geq 1$
$\delta_2(x)$	if $x \leq 1$	if $x \geq 2$
$\delta_3(x)$	if $x \leq 2$	if $x \geq 3$
$\delta_4(x)$	if $x \leq 3$	if $x \geq 4$
$\delta_5(x)$	if $x \leq 4$	if $x \geq 5$
$\delta_6(x)$	if $x \leq 5$	if $x \geq 6$
$\delta_7(x)$	if $x \leq 6$	if $x \geq 7$
$\delta_8(x)$	if $x \leq 7$	if $x \geq 8$
$\delta_9(x)$	Always	Never

Our problem, then, will be to select from these 10 reasonable decisions, that one which will minimize the expected or Bayes risk. Substituting into Equation (2.8a) and (2.8b), we indicate the risk of each (δ_s, θ_j) combination by:

$$R(\delta_s, \theta_1) = 2 \sum_{x=s}^{8} \binom{8}{x} .2^x .8^{8-x}$$

and

$$R(\delta_s, \theta_2) = 1 - \sum_{x=s}^{8} \binom{8}{x} .8^x .2^{8-x} .$$

Thus, from Equation (2.9), we have the Bayes risk given by

$$r(\delta_s, \rho) = .6 + \sum_{x=s}^{8} \left[.8\binom{8}{x} .2^x .8^{8-x} - .6\binom{8}{x} .8^x .2^{8-x} \right] . \quad (2.12)$$

Our theory tells us that this function is a minimum if we take s to be the next integer greater than x_o, where x_o is given by:

$$x_o = \frac{\ell n(2) + \ell n(2/3) + 8\ \ell n(4)}{\ell n(4) + \ell n(4)}$$

$$= 4.104 .$$

And, therefore, the Bayes risk $r(\delta_s, \rho)$ will be a minimum if we choose $\delta_5(x)$. Substituting $s = 5$ into Equation (2.12), we see that the minimum Bayes risk is $r(\delta_5, \rho) = .042$. To convince the still skeptical reader that δ_5 does, in fact, lead one to the minimum Bayes risk, in Table 16 we have exhibited $r(\delta_s, \rho)$ for each of our decision rules.

Table 16

$r(\delta_s, \rho)$

for $c/b = 2$, $p = .6$, and $t = 8$

s	$r(\delta_s, \rho)$
0	.800
1	.666
2	.397
3	.163
4	.051
5	.042
6	.123
7	.298
8	.499
9	.600

As promised, we shall now reconsider the preceding problem using extensive form analysis. It will be recalled that analyses of the extensive form have two major advantages. First, it is unnecessary to exhibit all reasonable decision rules. And secondly, in minimizing the expected loss, it is necessary to consider only that value of x actually obtained. As we shall see, these simplifications will make this analysis almost trivial.

Recall that the likelihood of X given θ is given by

$$p(x|\theta) = \begin{cases} \binom{t}{x}\beta^x(1-\beta)^{t-x} & \text{for } \theta = \theta_1 \text{ (nonmaster)} \\ \binom{t}{x}\alpha^x(1-\alpha)^{t-x} & \text{for } \theta = \theta_2 \text{ (master)} \end{cases}$$

where $x = 0, 1, 2, \ldots, t$ [see Equation (2.6)], and that $\Pr(\theta = \theta_2) = p$. Combining these two probabilities we see that the joint probability density of θ and X is given by

$$p(x, \theta_j) = \begin{cases} \binom{t}{x}(1-p)\beta^x(1-\beta)^{t-x} & j = 1 \\ \binom{t}{x}p\alpha^x(1-\alpha)^{t-x} & j = 2 \end{cases}$$

where $x = 0, 1, 2, \ldots, t$. Therefore, by Bayes theorem, the posterior distribution of θ is given by

$$p(\theta_j|x) = \begin{cases} \dfrac{(1-p)\beta^x(1-\beta)^{t-x}}{(1-p)\beta^x(1-\beta)^{t-x} + p\alpha^x(1-\alpha)^{t-x}} & \text{for } j = 1 \\ \dfrac{p\alpha^x(1-\alpha)^{t-x}}{(1-p)\beta^x(1-\beta)^{t-x} + p\alpha^x(1-\alpha)^{t-x}} & \text{for } j = 2 \end{cases} . \qquad (2.13)$$

Combining Table 14 with Equation (2.13), we see that the expected posterior loss under d_1 is given by

$$\mathcal{E}_{\theta|x}[L(d_1, \theta)] = \frac{bp\alpha^x(1-\alpha)^{t-x}}{(1-p)\beta^x(1-\beta)^{t-x} + p\alpha^x(1-\alpha)^{t-x}}$$

and that under d_2 by

$$\mathcal{E}_{\theta|x}[L(d_2, \theta)] = \frac{c(1-p)\beta^x(1-\beta)^{t-x}}{(1-p)\beta^x(1-\beta)^{t-x} + p\alpha^x(1-\alpha)^{t-x}} .$$

Therefore, extensive form analysis leads one to the decision rule given by

$$\delta^*(x) = \begin{cases} d_1 & \text{if } bp\alpha^x(1-\alpha)^{t-x} < c(1-p)\beta^x(1-\beta)^{t-x} \\ d_2 & \text{if } bp\alpha^x(1-\alpha)^{t-x} > c(1-p)\beta^x(1-\beta)^{t-x} \end{cases}$$

or equivalently,

$$\delta^*(x) = \begin{cases} d_1 & \text{if } x \ln\left[\frac{\alpha(1-\beta)}{\beta(1-\alpha)}\right] < \ln\left(\frac{c}{b}\right) + \ln\left(\frac{1-p}{p}\right) + t \ln\left(\frac{1-\beta}{1-\alpha}\right) \\ d_2 & \text{if } x \ln\left[\frac{\alpha(1-\beta)}{\beta(1-\alpha)}\right] > \ln\left(\frac{c}{b}\right) + \ln\left(\frac{1-p}{p}\right) + t \ln\left(\frac{1-\beta}{1-\alpha}\right) . \end{cases}$$

By the argument following Equation (2.11), we see that <u>when $\alpha > \beta$</u> this decision rule is equivalent to that reached under analysis of the normal form. Thus, when $\delta > \beta$,

$$\delta^*(x) = \begin{cases} d_1 & \text{if } x < x_o \\ d_2 & \text{if } x > x_o \end{cases}$$

where

$$x_o = \frac{\ln(c/b) + \ln[(1 - p)/p] + t \ln[(1 - \beta)/(1 - \alpha)]}{\ln[(1 - \beta)/(1 - \alpha)] + \ln[\alpha/\beta]} \quad .$$

Consider again this last experimental situation. A moment's reflection will confirm that increases in the number of tasks in the test will increase the overall loss in the decision problem. Each task included in the test, for example, increases the time the student must devote to testing and involves some commitment of facilities and, perhaps, time on the part of a teacher. If we can assume that a certain fixed cost k is incurred for each task in the test, we can reformulate our decision problem into two parts to take account of this additional loss: The first being the selection of the critical number of items I to be answered correctly, and the second being the selection of the optimal number of tasks T to minimize the total Bayes risk $r[\delta(t)^*, \rho] + tk$ [where $\delta(t)^*$ denotes the optimum decision rule when the test contains t items]. The final term, tk, may be thought of as the cost of testing.

Within the framework of normal form analysis, we would seek a cutting score (s) and a number of test tasks (t) that would minimize the total Bayes risk

$$bp \sum_{x=0}^{s-1} \binom{t}{x} \alpha^x (1 - \alpha)^{t - x} + c(1 - p) \sum_{x=s}^{t} \binom{t}{x} \beta^x (1 - \beta)^{t - x} + tk \quad .$$

The constant unit cost or unit loss k must be on the same scale as the original loss function, if we are to have a valid total Bayes risk.

For each fixed value of t, we already know how to select the critical number of items I to be answered correctly. Consequently, with the aid of a computer, it is easy to determine values of $r[\delta(t)^*, \rho] + tk$ for a range of test lengths t . We can then search the display of values of the

total Bayes risk, searching for T, the optimum level of t . An analogous approach applies within the framework of extensive form analysis.

In our example, with c = 2, b = 1, α = .8, β = .2, and p = .6, suppose the loss associated with administering one item is k = .01. This could happen, for example, if we think of each test task as taking up .01 as much time as an instructional unit. If the loss constant c = 2 was selected because it is associated with the time loss that will result in repeating two instructional units and b = 1 because it is associated with the time loss that will result in repeating one unit, then k will be on the same scale. From Table 17, we can see something of the shape of the total Bayes risk function for this example. As the table clearly indicates, in this case the decision maker would choose a test length of t = 7 and a critical test score I = 4.

Table 17

Display of Total Bayes Risk

for Various Values of t

No. of Items (t)	Critical Score (I)	$r[\delta(t)^*, \rho] + .01\ t$
0	1	.6000
1	1	.2900
2	2	.2680
3	2	.1756
4	3	.1702
5	3	.1311
6	4	.1329
7	4	.1167
8	5	.1221
9	5	.1174
10	6	.1248
11	6	.1263
12	7	.1348
13	7	.1398
14	8	.1489
15	8	.1559
16	9	.1654
17	9	.1736
18	10	.1833
19	10	.1922
20	11	.2020

$b = 1$; $c = 2$; $\alpha = .8$; $\beta = .2$; $p = .6$; $k = .01$

3. Utility Theory

In this section, we turn our attention to a theory which is much stronger than that illustrated in the previous examples. We will now abandon the notion of loss which we relied upon so heavily in Section 2, in favor of the more generally applicable notion of utility. The major difficulty with using loss functions as previously described, lies in the fact that we have simply assumed their existence. Because of the apparent reliance of the loss function on some scale, be it economic, social, political, or other, it is by no means obvious that such a function should exist.

In marked contrast, utility notions do not require that we invent a different and in some sense arbitrary scaling procedure for each problem we meet. Instead, utility theory uses the notions of ordered personal preferences or desirability of outcomes to scale the consequences of each (d, θ) pair. Although several axiom systems have been proposed to insure the existence of utility functions, these axioms generally require only the very basic relationships between preferences which rationality demands. Although these axioms contain many structural details concerning the nature of outcomes, preferences, and rewards, the most important characteristics of these axioms for applications seem to be the requirements for the comparability of any two outcomes and the coherence of the set of possible comparisons.[3] The first requirement merely assures that for any two outcomes A and B, precisely one of the following situations must obtain:

1) A is preferred to B, or

2) B is preferred to A, or

---

[3] The interested and mathematically able reader is referred to De Groot (1970, Chapter 7) for a detailed consideration of these axioms.

3) A and B are equally desirable outcomes and, we are, therefore, <u>indifferent</u> as to which occurs.

The second requirement is one of transitivity. It merely asserts that for any three outcomes A, B, and C, if A is preferred to B and B is preferred to C, then it must be that A is preferred to C. The point is that these requirements of comparability and coherence are both simple and reasonable. They are incorporated into our system of preferences without question, for it is generally agreed that any violation of these axioms in practice, if exposed, would be deemed ridiculous and one's system of preferences reconsidered.

As an aside, we note that those readers who like the notion of loss described in the previous section need not despair. As Lindley (1972) points out, in applications it typically seems true that one can define a suitable loss function by

$$L(d, \theta) = \underset{d}{\text{Max}}\{u(d, \theta)\} - u(d, \theta) \qquad (3.1)$$

if the number of outcomes is finite.

Although we will not devote a great deal of space to the problem of assessing one's utility, we will consider one method which will work in problems where there are a finite number of outcomes.

We can represent the set of outcomes by the following table:

	θ_1	θ_2	θ_3	. . .	θ_n
d_1	C_{11}	C_{12}	C_{13}	. . .	C_{1n}
d_2	C_{21}	C_{22}	C_{23}	. . .	C_{2n}
.	.	.	.		.
.	.	.	.		.
.	.	.	.		.
d_m	C_{m1}	C_{m2}	C_{m3}	. . .	C_{mn}

Where C_{22}, say, is the outcome or consequence associated with making decision d_2 when θ_2 is the true state of nature. Let C^* be the most preferred outcome and C_* be the least preferred in the table, and assume that you are given a lottery ticket with a v percent chance of "winning" C^* and a 100 - v percent chance of "winning" C_*. Further, assume that someone has offered to take the ticket off your hands in exchange for C_{ij}. Your task is to discover that value of v such that you would be willing to flip a fair coin to decide between the alternatives:

1) A ticket with a v percent chance on C^* and a 100 - v percent chance on C_*, or

2) Selling your ticket for outcome C_{ij}.

The utility of outcome C_{ij} can then be defined by $u(C_{ij}) = v/100$. This procedure can then be followed for each C_{ij} in turn until utilities have been coherently assigned to each of the outcomes.

We now return to our initial example to illustrate this procedure. Recalling that we have two states, nonmaster and master, and two reasonable decisions, retain and advance, we summarize the outcomes in Table 18.

Table 18

	possible outcomes C_{ij}	
	θ_1 (nonmaster)	θ_2 (master)
d_1 (retain)	C_{11}	C_{12}
d_2 (advance)	C_{21}	C_{22}

Surely the most desirable outcomes are C_{11} and C_{22}. In either of these cases, we correctly classify the student, so it is probably unreasonable to believe that one should be preferred to the other. Furthermore, on the presumption that if a nonmaster is advanced, he will not only lose the

time required to complete the next unit, but may also become frustrated and discouraged, let us assume that misclassifying a master is much more desirable than misclassifying a nonmaster. So in terms of our notation, we have $C_{11} = C_{22} = C^*$ and $C_{21} = C_*$. In this simplified example, the only remaining problem is to determine $u(C_{12})$. What we need to determine is that value v such that our decision maker would be willing to flip a fair coin to decide which gamble he will take:

1) A lottery which pays off C^* (a correctly classified student) v percent of the time and C_* (a misclassified nonmaster) (100 - v) percent of the time; or

2) A sure C_{12} (misclassified master).

Admittedly, specifying v is not an easy task, but it can be done. In order to accomplish this, our decision maker might be aided by considering "how much better" or more desirable correctly classifying a student is than misclassifying a nonmaster and compare this with how much better correctly classifying a student is than misclassifying a master. If, for example, correctly classifying a student gives you 10 "utiles" more than misclassifying a nonmaster and only 5 "utiles" more than misclassifying a master, then v for C_{12} would be 50 percent. This says that misclassifying a master is half-way between misclassifying a nonmaster and correctly classifying a student on a "utiles" or desirability scale. Assuming that, in fact, v = 50 then $u(C_{12}) = v/100 = .5$. Carrying out the above procedure for C_{11}, C_{21}, and C_{22}, we see that it must be true that $u(C_{11}) = u(C_{22}) = 1$ and $u(C_{21}) = 0$. Summarizing this in Table 19, we have

Table 19

Utilities of (d_i, θ_j)

	θ_1 (nonmaster)	θ_2 (master)
d_1 (retain)	1	.5
d_2 (advance)	0	1

In the next section, we will examine some classes or families of utility functions which may be used to describe a decision maker's preferences in the dichotomous or two-action decision problem. Analogous to the minimization of expected loss in extensive form analysis, decision theory with a utility function requires us to select that decision which will maximize the posterior expected or average utility. That is, we seek that $d_i (i = 1, 2)$ such that

$$\int u(d_i, \theta) p(\theta|x) d\theta$$

is a maximum. As usual, θ is the parameter which summarizes the state of nature and $p(\theta|x)$ is its posterior density.

In what follows, it will be important to recognize that if

$$\mathcal{E}_{\theta|x}[u(d_1, \theta)] > \mathcal{E}_{\theta|x}[u(d_2, \theta)] ,$$

then for $b > 0$

$$\mathcal{E}_{\theta|x}[bu(d_1, \theta) + c] > \mathcal{E}_{\theta|x}[bu(d_2, \theta) + c] .$$

That is, if d_1 is preferred to d_2 using the utility function $u(d_i, \theta)$, then d_1 will still be preferred using any positive linear transformation of $u(d_i, \theta)$. A similar demonstration is valid when the direction of

the inequalities is reversed. So positive linear transformations of utility functions will not alter the ultimate decision. This means that for decision purposes, a utility function needs to be determined only up to a positive multiplicative constant and an additive constant.

4. Utility and the Two Action Problem

As we saw in the last section, utility functions are imprecise things. In practical situations, we are often unable to specify our utility associated with various (d_i, θ) pairs with anything approaching mathematical precision. Generally, the best we can do is find some approximation which agrees fairly well with our subjective evaluation of the payoffs.

Our purpose in this section will be to illustrate a variety of families or sets of utility functions which have proven useful in applications. In selecting families of utility functions for inclusion in this section, we have used two principal criteria. First of all, we have sought to include families which are mathematically tractable in the sense that their expected values are easily calculated for standard distributions. Secondly, we have sought families which permit an acceptable compromise between having too many parameters for the decision maker to conveniently specify, and being so restricted that significant aspects of the decision maker's preferences cannot be expressed.

4.1 Threshold Utility

With threshold utility, like threshold loss discussed in Section 2, we separate (or partition) the possible values of our state parameter θ into a number of mutually exclusive subsets. For continuous θ we might consider the partition $\{A_1, A_2\}$ where $A_1 = \{\theta | \theta < \theta_o\}$ and $A_2 = \{\theta | \theta \geq \theta_o\}$ for some θ_o. Thus, for decision purposes, those values of θ which are less than some point, θ_o, will be considered as a set and denoted A_1. Those values of θ which are greater than or equal to θ_o will be grouped together as A_2. This partition would be analogous to our previous

examples where θ was a measure of student ability and θ_o was that point which separated the masters from the nonmasters on an educational test.

Another possibility with threshold utility would be to use a slightly finer partition of θ . This might be accomplished, for example, by preselecting three points instead of one: θ_1, θ_2, and θ_3, say. We could then use the partition $\{B_1, B_2, B_3, B_4\}$ where $B_1 = \{\theta|\theta < \theta_1\}$; $B_2 = \{\theta|\theta_1 \leq \theta < \theta_2\}$; $B_3 = \{\theta|\theta_2 \leq \theta < \theta_3\}$; and $B_4 = \{\theta|\theta \geq \theta_3\}$. Such a scheme might be useful if we needed to differentiate those masters who "just barely made the grade" from those who had truly assimilated the material, with analogous distinctions for the nonmasters. Naturally, even finer partitions could be used if the situation warranted it.

Returning to our dichotomous partition $\{A_1, A_2\}$, let us consider it in more detail. If we denote the utility associated with each (d_i, A_j) pair by $u(d_i, A_j)$, then in the two action problem, we may represent the threshold utility function as in Table 20.

Table 20

$u(d_i, A_j)$

	$A_1(\theta < \theta_o)$	$A_2(\theta \geq \theta_o)$
d_1	a	b
d_2	c	d

Decision or action d_1 is then to be preferred whenever

$$\mathcal{E}_{\theta|x}[u(d_1, A_j)] > \mathcal{E}_{\theta|x}[u(d_2, A_j)] ,$$

or whenever

$$a\{Pr[\theta < \theta_o]\} + b\{Pr[\theta \geq \theta_o]\} > c\{Pr[\theta < \theta_o]\} + d\{Pr[\theta \geq \theta_o]\} . \quad (4.1)$$

In decision problems similar to those discussed earlier in the context of educational testing, the d's involve decisions about the true state of nature θ: "Is he a master or a nonmaster?" When this is the case, it will usually happen that either a and d or b and c will be associated with "correct" decisions, and therefore, ought to be larger than the other utilities in their respective columns in the table. Since the labeling of d_1 and d_2 is arbitrary, we will assume that a and d are the correct decisions in this analysis. Furthermore, since we demonstrated at the end of the last section that rescaling utility by a positive linear transformation does not affect our decisions, we will assume that all utilities in the table lie in the interval between zero and one. Applying these stipulations to Equation (4.1), we see that decision d_1 will be preferred whenever

$$(a - c)\{Pr[\theta < \theta_o]\} > (d - b)\{Pr[\theta \geq \theta_o]\}$$

or alternatively,

$$\frac{d - b}{a - c} < \frac{Pr[\theta < \theta_o]}{Pr[\theta \geq \theta_o]} \quad . \tag{4.2}$$

Since

$$Pr[\theta \geq \theta_o] = 1 - Pr[\theta < \theta_o] \ ,$$

we have

$$\frac{d - b}{(a - c) + (d - b)} < Pr[\theta < \theta_o] \ . \tag{4.3}$$

By a similar argument, d_2 will be preferred whenever

$$\frac{d - b}{(a - c) + (d - b)} > Pr[\theta < \theta_o] \ .$$

The reader should note the similarity between Equation (4.3) and Equation (2.2). In fact, if $a = d = 1$ and one considers the loss as specified in Equation (3.1), then it is clear that the equations are identical. And so by simply rephrasing the example of Section 2.3 in terms of utility instead of loss, we see that the posterior cutting score is given by

$$\frac{1 - b}{2 - b - c} = \Pr[\theta < \theta_o | x_o] .$$

Consider the following numerical example previously discussed in Section 2.3. Suppose that the test score X has a normal distribution with unknown mean θ and unknown variance ϕ . Further suppose that after careful consideration of all collateral information available, our decision maker is able to adequately summarize his prior beliefs about θ and ϕ for the student under study in terms of a gamma-normal distribution with parameters $m = 9$, $w_o = 80$, and $R^2 = 144$. By Equation (2.4), we see that

$$t(\theta|x) = \frac{[9(10)]^{\frac{1}{2}}[\theta - (9*80 + x)/10]}{[144 + 9(x - 80)^2/10]^{\frac{1}{2}}} \qquad (4.4)$$

has a student t distribution on 9 degrees of freedom.

In order to apply Equation (4.3), we need two additional pieces of information: A critical true score θ_o, and a utility function $u(d_i, \theta)$. As before we let $\theta_o = 75$. For purposes of this example, we describe our utility function by the following table.

$u(d_i, \theta)$

	$\theta < \theta_o$	$\theta \geq \theta_o$
d_1 (retain)	.9	.5
d_2 (advance)	0	1

Applying Equation (4.3), we see that we will retain the student if $Pr[\theta < \theta_o] > .5/(.9 + .5) \doteq .36$. And by Equation (4.4), this is equivalent to $Pr[t(\theta|x) < t(75|x)] > .36$, where $t(\theta|x)$ is distributed as a student's t with 9 degrees of freedom. If we have already collected our data so that x is known, this decision criterion can be applied directly. However, if x is not yet known, we can apply the same argument used in Section 2.3 to determine the now familiar cutting score, x_o, which divides the observation scale into two disjoint decision regions. There are two steps in the determination of x_o. First, we must locate the point t_o in a table of the central t distribution with 9 degrees of freedom, which has percentile rank equal to 36. Secondly, we must solve the equation

$$t_o = -.37 = \frac{(90)^{\frac{1}{2}}[75 - (720 + x_o)/10]}{[144 + 9(x_o - 80)^2/10]^{\frac{1}{2}}} \qquad (4.5)$$

for x_o. Working out the algebra, we find $x_o = 43.5$ and $x_o = .635$. Substituting these two possible solutions into Equation (4.5), we see that only $x_o = 43.5$ satisfies the equation as stated. Thus, whenever the observed x is greater than 43.5, extensive form analysis will lead the decision maker to advance the student. Of course, whenever x is less than 43.5, the student will be retained.

We turn now to another class of utility functions. This time, we will treat utility as a continuous function of θ rather than a discrete one. Continuous utility might be considered as the limiting case of threshold utility as the partition grows increasingly fine.

4.2 Linear Utility

We direct our attention first to the simplest of continuous utility functions, that is, those which are linear in θ . In the linear two-action case, we define utility by functions of the following form.

$$u(d_i, \theta) = \begin{cases} e + f\theta & i = 1 \\ g + h\theta & i = 2 \end{cases} \tag{4.6}$$

The reader should note that what we have done here is define utility as a separate linear function for each possible decision, d_i . Thus, if decision one is chosen, the payoff or utility is to be a linear function of the state parameter θ with slope f and intercept e . For decision d_2, the slope is h and the intercept is g .

The existence of a breakeven or indifference value θ_o of the state parameter θ imposes the condition that

$$e + f\,\theta_o = g + h\,\theta_o, \text{ or } \theta_o = (e - g)/(h - f) \quad .$$

In our attempt to maximize expected utility, we will select action d_1 if

$$\mathcal{E}[(e + f\theta)|x] > \mathcal{E}[(g + h\theta)|x] \quad .$$

If $\mu_\theta = \mu_{\theta|x}$ denotes the posterior mean of θ, this implies that action d_1 is taken if

$$e + f\,\mu_\theta > g + h\,\mu_\theta \quad .$$

In other words, with linear utility, the action taken depends only upon the mean of the posterior distribution of the state parameter θ, other attributes of the distribution are irrelevant for decision purposes.

If we index our decisions so that $h - f > 0$, we take action d_1 whenever $\mu_\theta < \theta_o$ and action d_2 whenever $\mu_\theta > \theta_o$. Figure 4.1 illustrates this

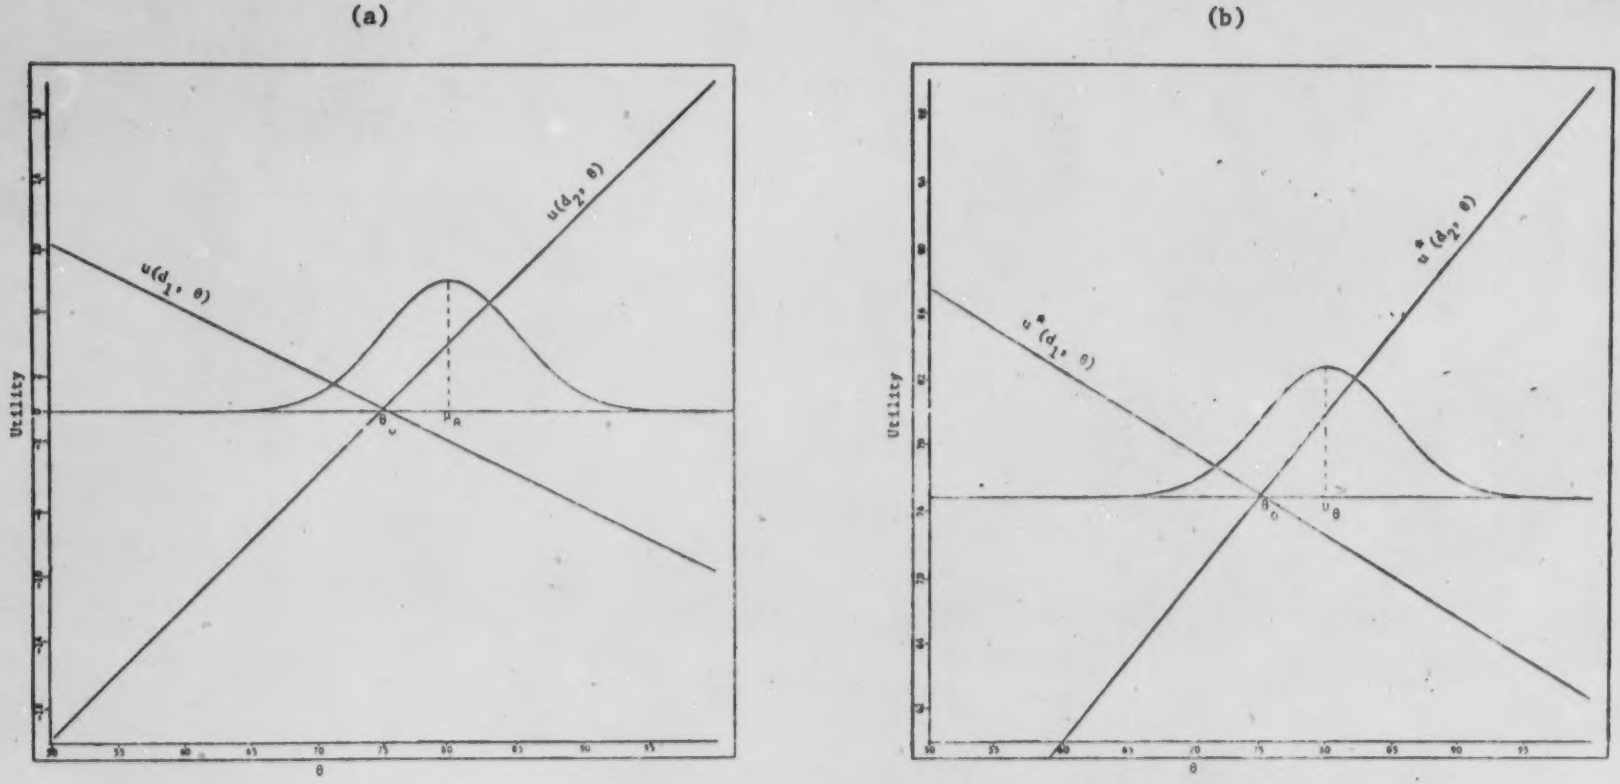

Figure 4.1

Linear Utility

Figure (a) above illustrates linear utility of the form of Equation (4.6) with constants e = 30.2, f = -.4, g = -59, and h = .8. Reparameterizing (a) according to $u^*(d_1, \theta)$, Equation (4.7), the utility of θ_o is equated to θ_o and the axes are rotated so that the d_2 branch has a slope equal to one. This reparameterized form of (a) is illustrated in (b). In this illustration, since $u^*(d_2, \mu_\theta) > u^*(d_1, \mu_\theta)$, extensive form analysis will lead the decision maker to select action d_2.

situation graphically. When linear utility is used, one needs to calculate only the utilities of each decision at the posterior mean μ_θ . The decision with the highest utility at μ_θ should then be selected.

As the general linear utility function now stands, Equation (4.6), we need to determine the four constants e, f, g, and h before it is completely specified. However, if we employ the flexibility afforded by the requirement that a utility function needs to be determined only up to a positive linear transformation, we can reduce the number of unknown constants to two. Thus, if $h > 0$ in Equation (4.6), we may redefine $u(d_i, \theta)$ by making the following positive linear transformation

$$u^*(d_i, \theta) = [u(d_i, \theta) - g]/h \ .$$

And so

$$u^*(d_i, \theta) = \begin{cases} e' + f'\theta & i = 1 \\ \theta & i = 2 \end{cases} \qquad (4.7)$$

where $e' = (e - g)/h$ and $f' = f/h$.

The nature of our assumption that $h > 0$ for this transformation to be valid cannot be overemphasized. The condition that $h > 0$ is equivalent to the statement that for decision d_2, utility is a strictly increasing function of the state parameter θ . In terms of our previous examples where we considered θ to be an ability index, $h > 0$ would make sense only if d_2 were the decision to advance the student. For if d_2 were the decision to retain him at the present level, we would be in the untenable position of asserting that as ability increases, the utility or desirability of retaining the student at the present level also increases.

If one is careful about making such transformations, this limitation will not cause serious problems. In applications, it is usually the case that the utility or desirability of one of the decisions will increase with the state parameter θ . Thus, all one needs to do is label that decision d_2 and result (4.7) is completely general.

We turn now to an illustration of one of the most direct methods available for determining the constants e' and f' of Equation (4.7). In order to make this method work, the decision maker must be able to specify two ordered pairs (θ_1, θ_2) and (θ_1', θ_2') such that

$$u(d_1, \theta_1) = u(d_2, \theta_2)$$

and

$$u(d_1, \theta_1') = u(d_2, \theta_2') .$$

Substituting the equivalents of these expressions from Equation (4.7), we have

$$e' + f'\theta_1 = \theta_2$$

and

$$e' + f'\theta_1' = \theta_2' .$$

Solving this system of equations, we find that

$$f' = f/h = \frac{\theta_2 - \theta_2'}{\theta_1 - \theta_1'}$$

and

$$e' = \frac{e - g}{h} = \theta_2 - f'\theta_1 .$$

To illustrate the simplicity of using linear utility, consider the following example. After completing a unit of Individually Prescribed Instruction, a student is given a 16 item test to determine whether or not he has mastered the material. From considerable past experience, our decision maker knows that 50% of those students obtaining scores of 12 on the test are able to satisfactorily complete the next sequence. On the basis of this information, he feels that a "true score" of twelve is the minimum necessary for advancing the student to the next unit. Reparameterizing this true score in terms of proportion correct, we find that $\theta_o = .75$.

The next stage in our decision making process consists of determining the posterior distribution on the state parameter θ, where in this problem θ denotes the true proportion correct. Using the techniques described in Novick, Lewis, and Jackson (1973), in Lewis, Wang, and Novick (1973), and in Wang (1973), our decision maker is able to determine a posterior distribution on $\gamma = \sin^{-1}\sqrt{\theta}$. Although this distribution is rather complicated and apparently does not exist in closed form, its precise specification is actually irrelevant for the decision-making process when linear utility is used. Under linear utility, if we can determine or at least approximate the expected value of θ, we will have gleaned all the information from the posterior distribution necessary to make our decision.

Lewis, Wang, and Novick (1973) estimate μ_θ by transforming $\mathcal{E}(\gamma|x)$ according to the equation

$$\text{est of } \mu_\theta = \sin^2[\mathcal{E}(\gamma|x)] .$$

Since this estimate of μ_θ is, in fact, equal to the median of the posterior distribution of θ, it is likely to be a poor estimate only in those cases where the posterior distribution of θ is highly skewed. Furthermore, this is likely to be the case only when the true proportion correct is near either

zero or one. Thus, as long as the critical criterion score θ_o is not too close to either zero or one, errors in estimating μ_θ are unlikely to lead to incorrect advance or retain decisions of major consequence. It is true that if the true proportion correct, θ, is very close to θ_o, an incorrect decision is likely. However, linear utility implies that differences in the utility of the two decisions are not great, for points near the point θ_o. What this means for the decision making process is that with a linear utility function, the output of a readily available and easy to use computer program [see Lewis, Wang, and Novick (1973)] will enable our decision maker to determine a useful estimate of the posterior mean μ_θ at which to evaluate his utility function.

At the next stage in the process, our decision maker must specify his utility function. Actually in this example, very little needs to be done. For most reasonable linear utility functions, the utility associated with the decision to retain the student will have a smaller slope than that associated with the decision to advance the student. Since the two branches of $u(d_1, \theta)$ must intersect at θ_o, he will retain the student if $\mu_\theta < \theta_o$ and advance him if $\mu_\theta \geq \theta_o$. Thus, as long as our decision maker is <u>certain</u> that he will be satisfied with a linear utility function, in the dichotomous decision problem, all he really needs to determine is the ordinal relationship between f and h. If $f > h$ [i.e., the slope of $u(d_1, \theta)$ is greater than the slope of $u(d_2, \theta)$], he will select decision d_1 whenever $\mu_\theta \geq \theta_o$ and select d_2 whenever $\mu_\theta < \theta_o$. Of course if $f < h$, the situation is reversed.

The "catch" to the foregoing simplicity is that <u>the decision maker is usually not certain that he will be satisfied with a linear utility function until he tries to specify one</u>. In practice the utility function should be overspecified by indicating at least three pairs (θ_i, θ_j) such that

$u(d_i, \theta_i) = u(d_j, \theta_j)$. By overspecification, the decision maker is forced to carefully weigh the implications of a linear utility function.

In applications, linear utility functions seem to behave more reasonably in the neighborhood of the breakeven point, θ_o, than do threshold functions. In this region, the rewards and penalties for correct and incorrect decisions frequently change smoothly rather than abruptly.

The fact that linear utility functions are not bounded when θ is unbounded creates some problems. Several of the axiom systems that have been used to construct decision analysis require that utility functions be bounded. To this theoretical objection must be added the practical fact that unbounded utility functions simply cannot be interpreted far from the breakeven point. These objections are partially removed if the posterior probability distribution of the state parameter θ is fairly closely packed around the breakeven point. If there is almost no probability attached to extreme values, unbounded utility is of little practical importance.

4.3 Quadratic Utility

As we saw above, decisions involving linear utility functions depend only on the mean of the posterior distribution. Quadratic utility functions on the other hand, result in making decisions that depend on both the mean and the variance of the posterior distribution. We begin by defining quadratic utility by a function of the form:

$$u(d_i, \theta) = \begin{cases} -a(\theta - b)(\theta - c) & i = 1 \\ -e(\theta - f)(\theta - g) & i = 2 \end{cases} . \tag{4.8}$$

Observe that in order to use this utility function, we must specify the six constants: a, b, c, e, f, and g . The constants b, c, f, and g have

special meaning in the present parameterization of the utility function. These constants correspond to those values of θ where the utility of the respective decisions is zero. In the example which we considered extensively in Section 2, it might be argued that one reasonable and convenient point which might be used to fix the location of our utility function would be the indifference point θ_o . If we can use the permissible linear transformation of $u(d_i, \theta)$ to force the utility of θ_o equal to zero, we will have established a reference point upon which to judge the utilities associated with other values of θ . In fact, this task can be accomplished by defining

$$u^*(d_i, \theta) = u(d_i, \theta) + a(\theta_o - b)(\theta_o - c)$$

or equivalently

$$u^*(d_i, \theta) = u(d_i, \theta) + e(\theta_o - f)(\theta_o - g) \quad .$$

Recognizing that at the indifference point θ_o, $u(d_1, \theta_o) = u(d_2, \theta_o)$, we can rewrite $u^*(d_i, \theta)$ in the form

$$u^*(d_i, \theta) = \begin{cases} -a(\theta - \theta_o)(\theta - c') & i = 1 \\ -e(\theta - \theta_o)(\theta - g') & i = 2 \end{cases}$$

where $c' = b + c - \theta_o$ and $g' = f + g - \theta_o$. Since our permissible linear transformation allows us to specify a scale for utility as well as a location, we may reduce the number of constants to be specified even further by a transformation of the form

$$u^{**}(d_i, \theta) = u^*(d_i, \theta)/a$$

for $a > 0$. Thus,

$$u^{**}(d_i, \theta) = \begin{cases} -(\theta - \theta_o)(\theta - c') & i = 1 \\ -e'(\theta - \theta_o)(\theta - g') & i = 2 \end{cases}$$

where $c' = b + c - \theta_o$, $e' = \frac{e}{a}$, and $g' = f + g - \theta_o$. In order to determine the remaining constants, c', e', and g', the decision maker needs to specify other points on the two branches of the utility function. Although this could be done directly by actually specifying the utilities associated with (d, θ) pairs, we describe a method which is probably easier to use in most situations. This method will frequently work in situations where the utility of each decision seems to approach a maximum asymptotically as the deviation between θ and θ_o increases. In this case, it seems reasonable to situate our quadratic curves so that the convex side is up as illustrated in Figure 4.2. This is equivalent to specifying that both a and e in our original model, Equation (4.8), are positive. We can now fix two of the remaining constants by identifying θ_1 and θ_2 such that θ_1 is at the lower end and θ_2 is at the upper end of the feasible domain of θ. Since we have indicated that the utility approaches a maximum asymptotically as θ approaches these points, θ_1 and θ_2, it seems reasonable to require that the maximum on the d_1 branch occur at θ_1 and that the maximum on the d_2 branch occur at θ_2. This requirement is equivalent to the following system of equations:

$$\frac{d}{d\theta} u^{**}(d_1, \theta)\Big|_{\theta = \theta_1} = -(2\theta_1 - \theta_o - c') = 0$$

$$\frac{d}{d\theta} u^{**}(d_2, \theta)\Big|_{\theta = \theta_2} = -e'(2\theta_2 - \theta_o - g') = 0 \quad .$$

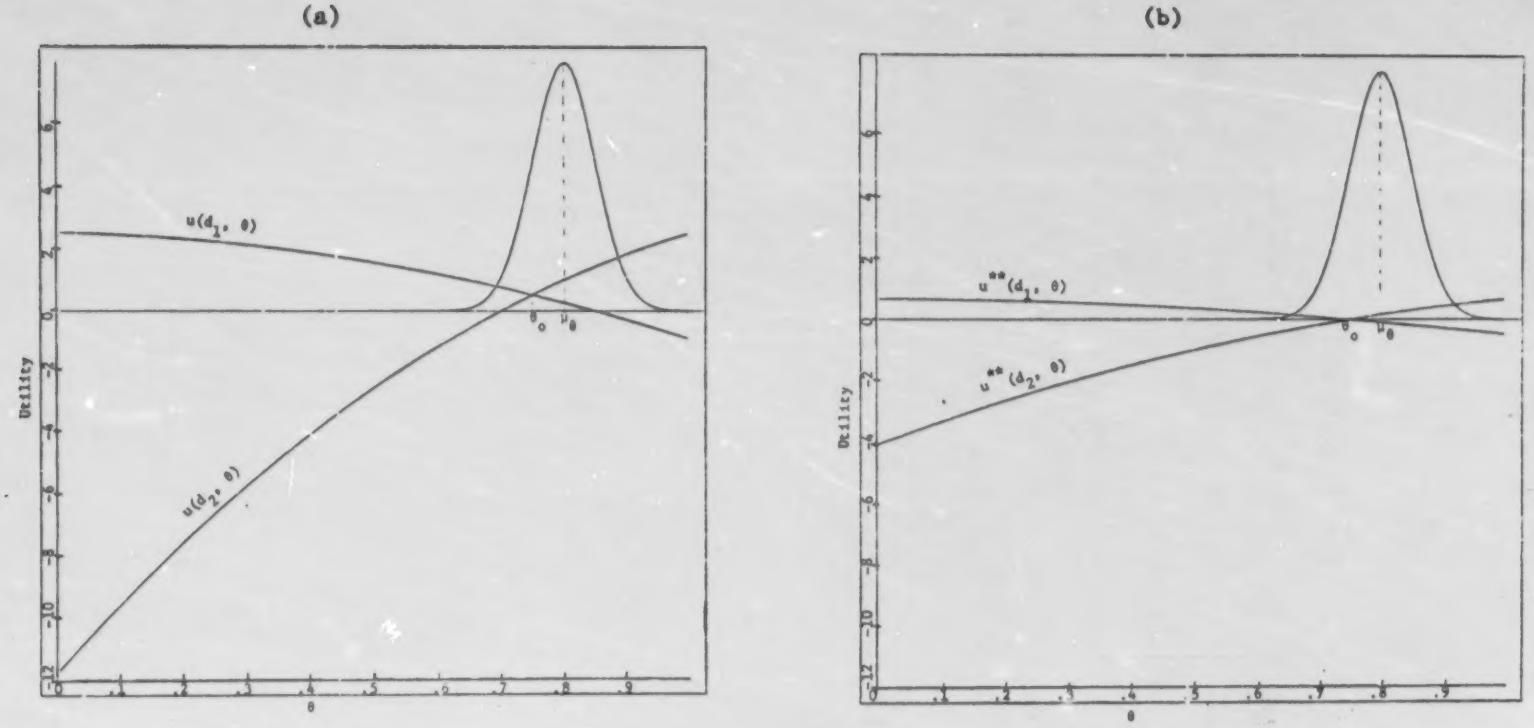

Figure 4.2

Quadratic Utility

Figure (a) illustrates a quadratic utility function of the form of Equation (4.8) with constants a = 3, b = -1, c = .85, e = 8.4, f = .7, and g = 2.0. Reparameterizing (a) according to $u^{**}(d_1, \theta)$ leads to (b). Reparameterization $u^{*}(d_1, \theta)$ changes the zero point on the utility scale so that the utility of θ_o is zero. This is equivalent to specifying b = f = θ_o, c = -.9, and g = 1.95 in Equation (4.8). Reparameterization $u^{**}(d_1, \theta)$ then alters the scale of utility so that a = 1, forcing e to equal 2.8.

Thus, $c' = 2\theta_1 - \theta_o$ and $g' = 2\theta_2 - \theta_o$. This procedure assumes, as is true in our examples, that large values of θ make action d_2 more desirable and that small values of θ make d_1 more favorable. Since the indexing of the decisions is arbitrary, however, this restriction is not serious. The remaining constant e' may now be determined by specifying a pair of state parameters (θ_3, θ_4) such that $u(d_1, \theta_3) = u(d_2, \theta_4)$. Since e' is the only unknown in this equation, it can be easily determined.

To illustrate these computations, let us reconsider the example used in the previous section with linear utility. In this example, the point at which the decision maker would be indifferent whether he advanced (d_2) or retained (d_1) the student was $\theta_o = .75$. Since θ is the "true" proportion correct, the minimum feasible θ is $\theta_1 = 0$ and the maximum is $\theta_2 = 1.0$. Thus, solving c' and g' in the equations above, we find $c' = -.75$ and $g' = 1.25$. If, in addition, the decision maker feels that $u(d_1, .7) = u(d_2, .85)$, say, then e' can be found by solving the equation

$$(.7 - .75)(.7 + .75) = e'(.85 - .75)(.85 - 1.25)$$

or

$$e' = 1.8 .$$

The utility function is illustrated in Figure 4.3. In general, the final decision will be for action d_1 if

$$\mathcal{E}[u^{**}(d_1, \theta)] > \mathcal{E}[u^{**}(d_2, \theta)]$$

where the expectation is taken with respect to the posterior distribution of θ . Thus, d_1 is to be preferred whenever

$$\sigma_\theta^2 + (\mu_\theta - \theta_o)(\mu_\theta - c') < e'\sigma_\theta^2 + e'(\mu_\theta - \theta_o)(\mu_\theta - g')$$

or

$$u^{**}(d_1, \mu_\theta) - \sigma_\theta^2 > u^{**}(d_2, \mu_\theta) - e'\sigma_\theta^2 ,$$

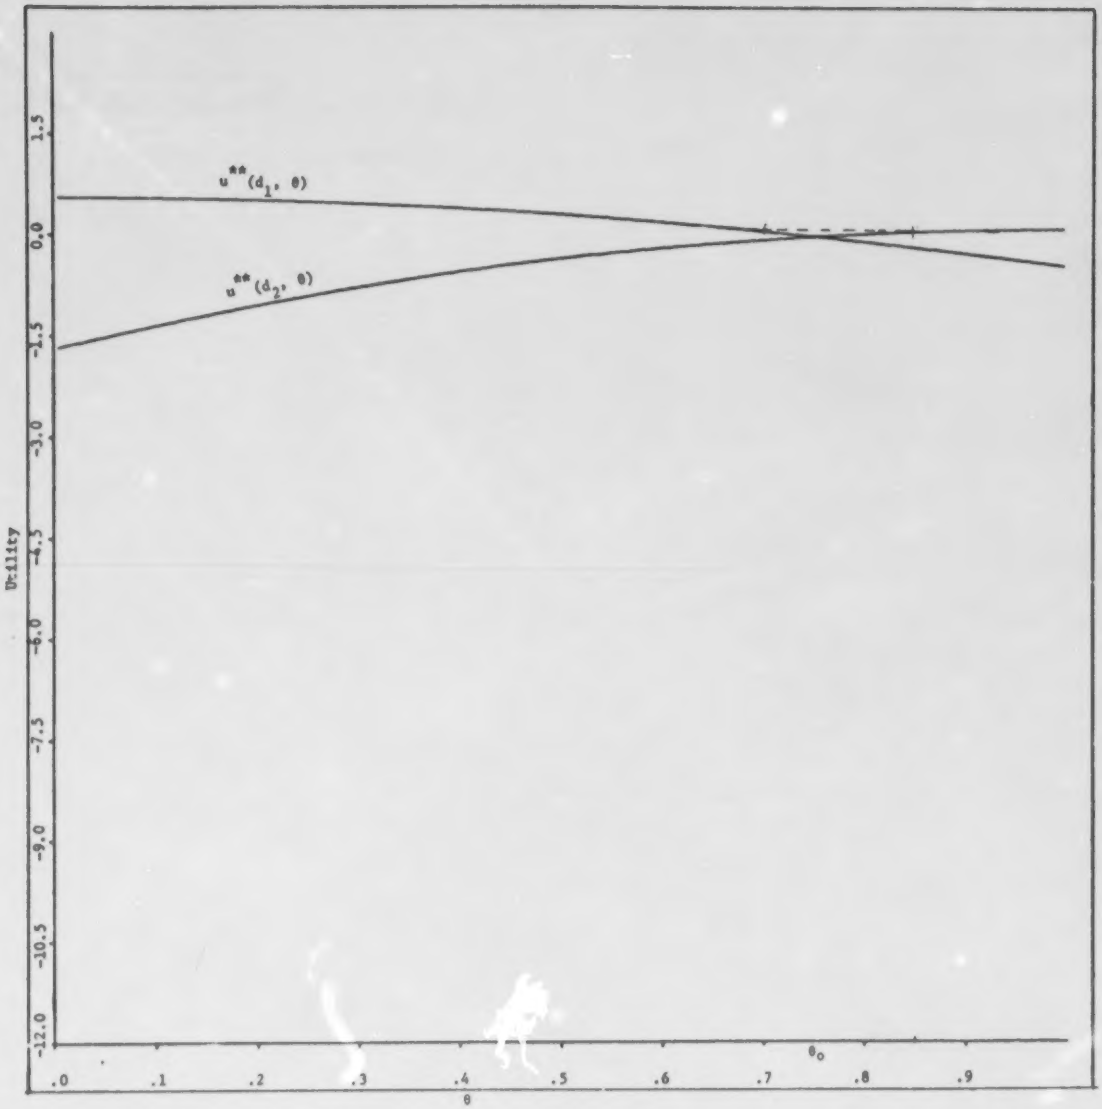

Figure 4.3

Quadratic Utility

This figure illustrates quadratic utility as transformed by $u^{**}(d_1, \theta)$. In this example, $c' = -.75$, $g' = 1.25$, $e' = 1.8$, $\theta_o = .75$, $\theta_1 = 0.0$, and $\theta_2 = 1.0$. Observe that $u^{**}(d_1, \theta)$ approaches its maximum at zero while $u^{**}(d_2, \theta)$ approaches its maximum at one. Also note that $u(d_1, .7) = u(d_2, .85)$.

where μ_θ and σ_θ^2 are the mean and variance of the posterior distribution of the state parameter θ . This result may be interpreted as calling for action d_1 whenever the utility of d_1 at μ_θ is greater than the utility of d_2 at μ_θ plus a correction factor. The correction factor depends upon the posterior variance of θ and consequently is a measure of the probable deviation of θ from μ_θ . For our Individually Prescribed Instruction examples, this relationship has important consequences. It indicates that for decision purposes, all we need from the posterior distribution of θ is its mean and its variance. In applications, we can often readily obtain these values, or approximations to them, even when the posterior distribution of θ does not exist in closed form.

4.4 Exponential Utility

Linear and quadratic utility functions have played important roles in applications of decision analysis. As we have seen, linear utility requires only that we evaluate the posterior mean, while quadratic utility requires both the posterior mean and variance. These simplifications of the decision process are extremely important, especially when the posterior distribution of θ is of a complicated form. Often we are able to estimate the mean and sometimes the variance of θ, even when the posterior density itself does not exist in closed form. As we shall see in this section, exponential utility also has a simplifying property which makes it particularly useful with many of the standard posterior density functions. Before illustrating this special property of exponential utility, we will exhibit the form of the function and perform our usual simplifications. We define exponential utility by a function of the following form:

$$u(d_i, \theta) = \begin{cases} c - a * \exp\{b\theta\} & i = 1 \\ c - f * \exp\{-g\theta\} & i = 2 \end{cases} \qquad (4.9a)$$

where the constants a, b, f, and g are positive. Notice that this is not the most general form available, since we require that a, b, f, and g be positive and that the same constant c appears under each decision rule. Although this simplification is made so that the estimates of the constants are more easily obtained, it has certain implications which the would be user should keep in mind. This particular formulation requires that for decision d_1, utility is a decreasing function of the state parameter θ. On the other hand, for decision d_2, utility must increase with increasing θ. To see this, all we need to do is rewrite $u(d_2, \theta)$ in the form $u(d_2, \theta) = c - f/\exp\{g\theta\}$. Clearly as θ increases, $f/\exp\{g\theta\}$ approaches zero. That is, as θ increases, the contribution of the second term to utility decreases rapidly, with $u(d_2, \theta)$ approaching c from below as an asymptote. A similar argument with respect to $u(d_1, \theta)$ shows that the utility of decision d_1 also approaches c from below, but this time with decreasing θ. This relationship is depicted graphically in Figure 4.4. Using our now familiar permissible linear transformation to eliminate some of the unknown constants, we let

$$u^*(d_i, \theta) = \frac{1}{a * \exp\{b\theta_o\}} \left[u(d_i, \theta) - (c - a * \exp\{b\theta_o\})\right] .$$

Since $a * \exp\{b\theta_o\} = f * \exp\{-g\theta_o\}$, we have

$$u^*(d_i, \theta) = \begin{cases} 1 - \exp\{b(\theta - \theta_o)\} & i = 1 \\ 1 - \exp\{-g(\theta - \theta_o)\} & i = 2 \end{cases} . \qquad (4.9b)$$

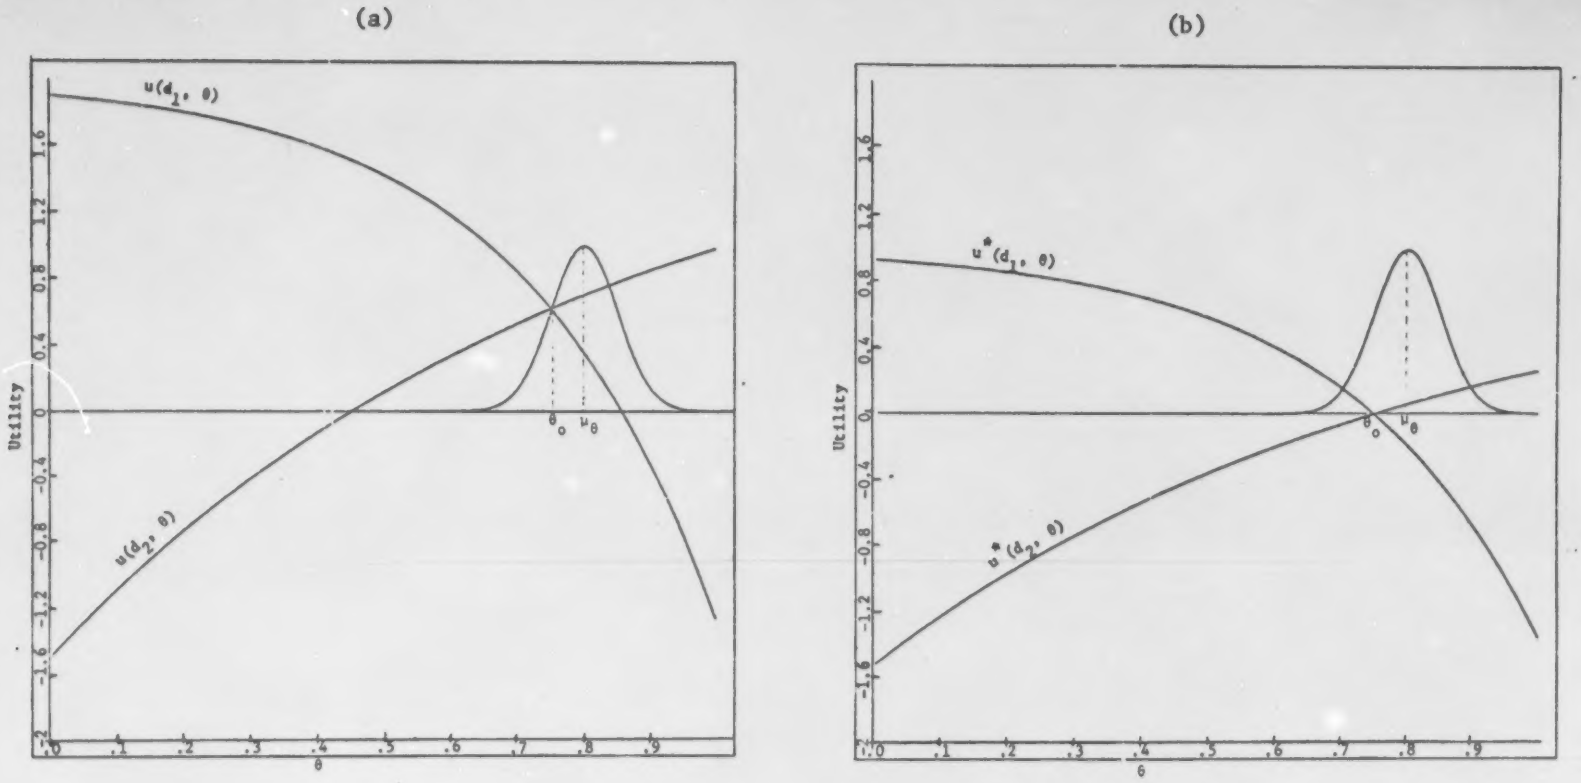

Figure 4.4

Exponential Utility

Figure (a) illustrates $u(d_1, \theta)$ from Equation (4.9) with a = 0.1, b = 3.5, c = 2.0, f = 3.5, and g = 1.24. Figure (b) is $u^*(d_1, \theta)$, the reparameterized form of figure (a). Upon reparameterization, the utility of θ_o is zero and the scale is changed so that c = a = f = 1.0. In its reparameterized form, $u^*(d_1, \theta)$, an exponential utility function is conpletely determined once θ_o and the slopes at θ_o, -b and g, of its two branches are specified.

When the decision maker turns to fixing the two parameters of the utility function, b and g, he must specify the precise utility of at least one point on each branch of the utility function in addition to θ_o . That is, he must specify a pair of points $\{(\theta_1, c_1), (\theta_2, c_2)\}$, such that $u(d_1, \theta_1) = c_1$ and $u(d_2, \theta_2) = c_2$. The reader is cautioned that this is not equivalent to what we have done in the past when we specified points (θ_1, θ_2) such that $u(d_1, \theta_1) = u(d_2, \theta_2)$. With exponential utility of the form of Equation (4.9), we must actually specify the values c_1 and c_2 (although they may, of course, be equal). Once this is done, the parameters b and g are completely determined, for we have

$$1 - \exp\{b(\theta_1 - \theta_o)\} = c_1$$

and

$$1 - \exp\{-g(\theta_2 - \theta_o)\} = c_2 .$$

Upon taking logarithms of both sides and solving, we have

$$b = \frac{\ell n\ (1 - c_1)}{(\theta_1 - \theta_o)} \tag{4.10}$$

and

$$g = -\frac{\ell n(1 - c_2)}{(\theta_2 - \theta_o)} \quad . \tag{4.11}$$

From these equations and the restriction on our model that b and g be positive, we see that if $\theta_1 < \theta_o$ and $\theta_2 > \theta_o$, then c_1 and c_2 must lie between zero and one. This is completely reasonable, however. In examining the transformed utility function, we see that it is zero at θ_o and increases to one as θ approaches either $-\infty$ or $+\infty$ for decisions d_1 and d_2, respectively (see Figure 4.4).

Probably the easiest way to specify c_1 in applications is to determine θ_1 such that the utility of (d_1, θ_1) is half way between the maximum possible utility under d_1 and the utility of (d_1, θ_o). By specifying a similar point θ_2 under d_2, the decision maker can check his coherence by observing whether $u(d_1, \theta_1) = u(d_2, \theta_2)$ subjectively. Having established his utility function, the decision maker will prefer decision d_1 whenever

$$\mathcal{E}[1 - \exp\{b(\theta - \theta_o)\} \mid x] > \mathcal{E}[1 - \exp\{-g(\theta - \theta_o)\} \mid x] .$$

But since expectation is a linear operator, this condition is equivalent to preferring d_1 whenever

$$1 - \mathcal{E}[\exp\{b(\theta - \theta_o)\} \mid x] > 1 - \mathcal{E}[\exp\{-g(\theta - \theta_o)\} \mid x]$$

or

$$\exp\{-b\theta_o\} \mathcal{E}[\exp\{b\theta\} \mid x] < \exp\{g\theta_o\} \mathcal{E}[\exp\{-g\theta\} \mid x]$$

or

$$\exp\{-(g + b)\theta_o\} \mathcal{E}[\exp\{b\theta\} \mid x] < \mathcal{E}[\exp\{-g\theta\} \mid x] .$$

But $\mathcal{E}[\exp\{t\theta\}] = M(t)$ has special significance for mathematical statisticians. In the statistical literature, $M(t)$ is referred to as the moment-generating function for θ. And because of the importance of these moment-generating functions, the integration necessary to evaluate the expected value has been worked out for most standard density functions. Therefore, if the posterior distribution of θ is one of the standard densities (Normal, Uniform, Triangular, Gamma, and others), the final stage in the decision-making process is merely a matter of "plugging in" the parameters of the posterior distribution on θ and those of the utility function. By reformulating the decision criterion in terms of moment-generating functions, we see that decision d_1 will be preferred whenever $\exp\{-(g + b)\theta_o\} M(b) < M(-g)$.

Returning to an example considered in Section 2.3, we assume that the posterior distribution of θ is normal with mean $\mu_\theta = (\tau\sigma^2 + x\phi)/(\sigma^2 + \phi)$ and variance $\sigma_\theta^2 = \sigma^2\phi/(\phi + \sigma^2)$. Since the moment-generating function for a normal variable is given by $M(t) = \exp\{t\mu_\theta + t^2\sigma_\theta^2/2\}$, decision d_1 will be preferred whenever

$$\exp\{b\mu_\theta + b^2\sigma_\theta^2/2 - (g + b)\theta_o\} < \exp\{-g\mu_\theta + g^2\sigma_\theta^2/2\}$$

or

$$\mu_\theta < \theta_o + (g - b)\sigma_\theta^2/2 \ . \qquad (4.12)$$

That is, the retain decision is preferred if the mean of the posterior distribution of θ is less than θ_o plus an adjustment which depends upon the variance of the posterior distribution and the relative utilities of the two decisions. Whether that adjustment is positive or negative, depends upon the sign of $g - b$. Since b and g are the magnitudes of the slopes of the utility function at θ_o for d_1 and d_2, respectively, the difference is a measure of the relative speed with which utility is changing on its two branches as θ moves away from θ_o . Thus, if $g > b$, the utility of d_2 is changing more rapidly in the vicinity of θ_o than the utility of d_1 . That is, when $g > b$, making a false, positive error will be relatively more expensive than an error of the false, negative variety for equal distances from θ_o . Consequently, the decision maker adjusts his critical point for μ_θ in a positive direction.

Let us reconsider a slight modification of the known variance numerical example presented in Section 2.3. As before, we assume that:

(1) $\rho(\theta) \sim N(80, 25)$ (prior on θ)

(2) $p(x|\theta) \sim N(\theta, 16)$ (likelihood or model density)

(3) $\theta_o = 75$ (critical criterion score).

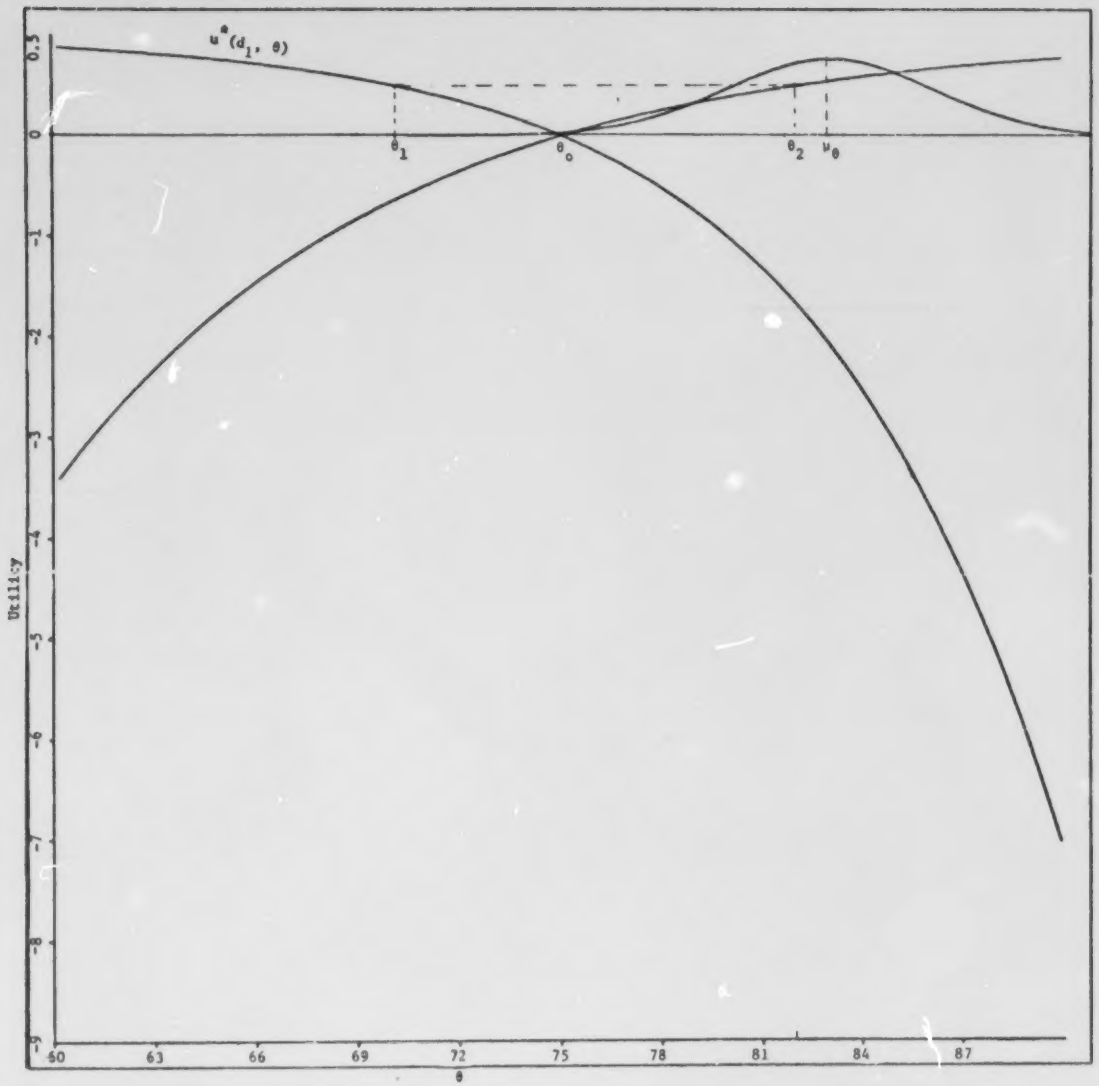

Figure 4.5

Exponential Utility

This figure illustrates exponential utility in the form of Equation (4.9b). In this example, we assumed that $\theta_1 = 70$ and $\theta_2 = 82$. Applying Equations (4.10) and (4.11), we concluded that b is equal to .14 and that g is equal to .10. The posterior distribution of θ is normal with mean equal to 83 and variance equal to 9.76.

We further assume that the student obtains a test score of $x = 85$. Thus, the posterior distribution of θ is normal with mean equal to 83 and variance equal to 9.76.

Next, our decision maker must specify his utility function. On the presumption that he would be satisfied with an exponential description like Equation (4.9), our decision maker would follow the procedures outlined in this section. His first task would be to subjectively determine θ_1 such that the utility of (d_1, θ_1) is half-way between $u(d_1, \theta_o)$ and the maximum possible utility on the d_1 branch. Next he must similarly determine θ_2 . Assuming that $\theta_1 = 70$ and $\theta_2 = 82$ and that these are coherent, Equations (4.10) and (4.11) indicate that $b = .14$ and $g = .10$. Applying Equation (4.12), we see that the student will be retained if $\mu_\theta < [75 + (-.08)9.76/2] = 74.61$. Since $\mu_\theta = 83$, the decision maker will certainly advance the student. This situation is illustrated in Figure 4.5.

4.5 Squared Exponential Utility

The final family of utility functions that we will consider for the two-action problem, will seem somewhat restricted in the amount of flexibility that it permits the decision maker. However, it is very compatible with normal posterior distributions and is frequently useful when other posteriors may be approximated by normal distributions. It is also a rather natural family of utility functions when the problem is an estimation problem and the act and parameter spaces coincide. The model for squared exponential utility is

$$u(d_i, \theta) = \begin{cases} 1 - \exp\left\{\dfrac{-a(\theta - \theta_o)^2}{2}\right\} & i = 1 \text{ and } \theta < \theta_o \\ 1 - \exp\left\{\dfrac{-b(\theta - \theta_o)^2}{2}\right\} & i = 2 \text{ and } \theta > \theta_o \\ 0 & \text{otherwise} \end{cases} \qquad (4.13)$$

82

where a and b are positive.

Turning to the specification of the parameters a and b, we follow the same procedure as in the case of exponential utility. That is, we specify the utility of at least one point other than θ_o for each branch of the utility function. If these points and their utilities are represented by the pairs $\{(\theta_1, c_1), (\theta_2, c_2)\}$, then

$$1 - \exp\{-\frac{a}{2}(\theta_1 - \theta_o)^2\} = c_1$$

and

$$1 - \exp\{-\frac{b}{2}(\theta_2 - \theta_o)^2\} = c_2 .$$

Upon taking logarithms and solving, we find that

$$a = -\frac{2\,\ell n(1 - c_1)}{(\theta_1 - \theta_o)^2}$$

and

$$b = -\frac{2\,\ell n(1 - c_2)}{(\theta_2 - \theta_o)^2} .$$

From these equations and the restriction on our model that a and b are both positive, we see that c_1 and c_2 must lie in the interval (0, 1). Examining Figure 4.6, we see that this is reasonable, for squared-exponential utility as described by Equation (4.13) is bounded between 0 and 1.

As with exponential utility, one way to determine (θ_1, c_1) and (θ_2, c_2) is to look for those points on the d_1 and d_2 branches such that

$$u(d_i, \theta_i) - u(d_i, \theta_o) = \max_{\theta}\{u(d_i, \theta)\} - u(d_i, \theta_i) .$$

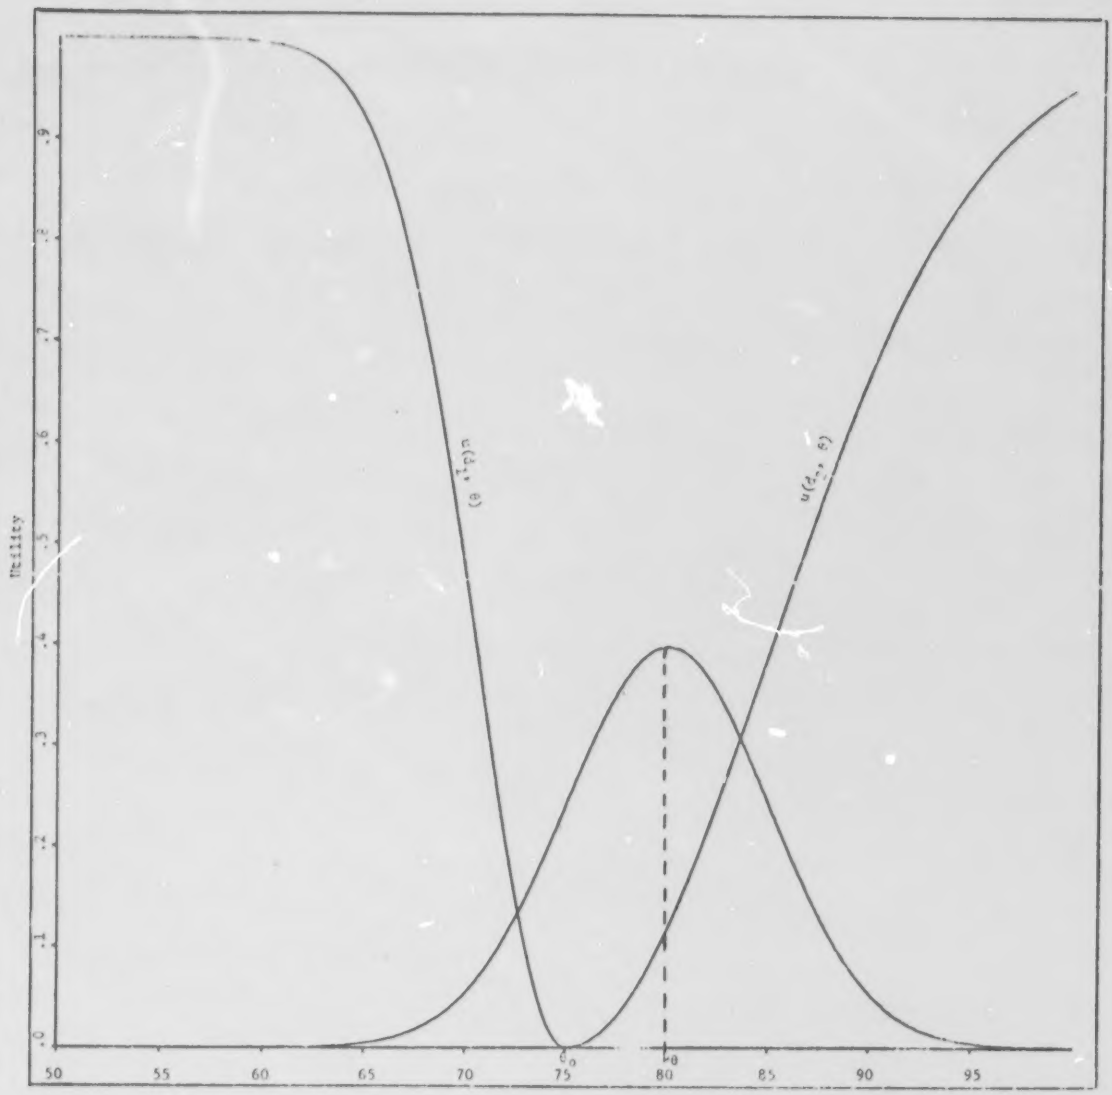

Figure 4.6

Squared Exponential Utility

This figure illustrates squared exponential utility in the form of Equation (4.13). In this figure, a equals .05 and b equals .01. Note that the utility function is bounded between zero and one, with $u(d_1, \theta)$ equal to zero for $\theta \geq \theta_o$ and $u(d_2, \theta)$ equal to zero for $\theta \leq \theta_o$.

That is, try to determine those points (d_i, θ_i) which have desirability approximately half-way between the desirability of (d_i, θ_o) and the maximum possible desirability under that decision. It will then be possible to verify the consistency (coherence) of these specifications by comparing $u(d_1, \theta_1)$ and $u(d_2, \theta_2)$, subjectively. If large differences are believed to exist in the payoff of these two situations, then some reconciliation will be necessary.

Action d_1 will be preferred if the expected or average posterior utility under d_1 is greater than that under d_2. That is, if

$$\int_{-\infty}^{+\infty} u(d_1, \theta)p(\theta|x)d\theta > \int_{-\infty}^{+\infty} u(d_2, \theta)p(\theta|x)d\theta \quad .$$

Since $u(d_1, \theta)$ and $u(d_2, \theta)$ are non-zero only in the regions $\theta < \theta_o$ and $\theta > \theta_o$, respectively, this inequality may be rewritten in the equivalent form

$$\int_{-\infty}^{\theta_o} u(d_1, \theta)p(\theta|x)d\theta > \int_{\theta_o}^{+\infty} u(d_2, \theta)p(\theta|x)d\theta \quad .$$

These integrals will be tractable if the posterior density of θ, $p(\theta|x)$, is normal with mean μ_θ and variance σ_θ^2. In this case,

$$u(d_1, \theta)p(\theta|x) = \frac{1}{\sqrt{2\pi\sigma_\theta^2}} \left[1 - \exp\left\{-\frac{a(\theta - \theta_o)^2}{2}\right\}\right]\left[\exp\left\{-\frac{(\theta - \mu_\theta)^2}{2\sigma_\theta^2}\right\}\right]$$

which may be rewritten in the form

$$u(d_1, \theta)p(\theta|x) = \frac{1}{\sqrt{2\pi\sigma_\theta^2}} \exp\left[-\frac{1}{2\sigma_\theta^2}(\theta - \mu_\theta)^2\right]$$

$$\cdot \frac{1}{\sqrt{2\pi\sigma_\theta^2}} \exp\left[-\frac{(a + \sigma_\theta^{-2})}{2}\left(\theta - \frac{a\theta_o + \sigma_\theta^{-2}\mu_\theta}{a + \sigma_\theta^{-2}}\right)^2\right] \exp\left[-\tfrac{1}{2}\left(\frac{a\sigma_\theta^{-2}}{a + \sigma_\theta^{-2}}\right)(\theta_o - \mu_\theta)^2\right] .$$

The first term, of course, is nothing but a normal density with mean μ_θ and variance σ_θ^2 . Ignoring that part of the second term which does not depend upon θ, we see that the second term is proportional to a normal distribution with mean $(a\theta_o + \sigma_\theta^{-2}\mu_\theta)/(a + \sigma_\theta^{-2})$ and variance $(a + \sigma_\theta^{-2})^{-1}$. Therefore, $\mathcal{E}[u(d_1, \theta)|x]$ may be written as

$$\mathcal{E}[u(d_1, \theta)|x] = \int_{-\infty}^{\theta_o} \frac{1}{\sqrt{2\pi\sigma_\theta^2}} \exp\left\{-\frac{(\theta - \mu_\theta)^2}{2\sigma_\theta^2}\right\} d\theta$$

$$\cdot \frac{1}{\sqrt{\sigma_\theta^2(\sigma_\theta^{-2} + a)}} \exp\left[-\tfrac{1}{2}\left(\frac{a\sigma_\theta^{-2}}{a + \sigma_\theta^{-2}}\right)(\theta_o - \mu_\theta)^2\right]$$

$$* \int_{-\infty}^{\theta_o} \frac{\sqrt{\sigma_\theta^{-2} + a}}{\sqrt{2\pi}} \exp\left[-\left(\frac{a + \sigma_\theta^{-2}}{2}\right)\left(\theta - \frac{a\theta_o + \sigma_\theta^{-2}\mu_\theta}{a + \sigma_\theta^{-2}}\right)^2\right] d\theta .$$

But each of these integrals is nothing but the probability in a tail area of a normal distribution. Therefore, we may express these integrals in terms of the percentile rank of θ_o . Thus, by standardizing θ in each integral so that each distribution is unit normal, we see that

$$100\,\mathcal{E}\,[u(d_1,\,\theta)\,|\,x] = \mathrm{PR}\left(\frac{\theta_o - \mu_\theta}{\sqrt{\sigma_\theta^2}}\right)$$

$$- \frac{1}{\sqrt{\sigma_\theta^2(\sigma_\theta^{-2} + a)}}\ \exp\left[-\tfrac{1}{2}\left(\frac{a\sigma_\theta^{-2}}{a + \sigma_\theta^{-2}}\right)(\theta_o - \mu_\theta)^2\right]\ \mathrm{PR}\left(\frac{\theta_o - \dfrac{a\theta_o + \sigma_\theta^{-2}\mu_\theta}{a + \sigma_\theta^{-2}}}{(a + \sigma_\theta^{-2})^{-\frac{1}{2}}}\right)$$

where PR(z) is the percentile rank of z with respect to the unit normal distribution.

Replacing a by b and reversing the limits of integration, we have a similar expression for the posterior expected utility under d_2. Therefore, our decision criterion may be written in the form: "Choose d_1 if

$$\mathrm{PR}(z_o) - \frac{\mathrm{PR}(z_1)}{\sqrt{\sigma_\theta^2(\sigma_\theta^{-2} + a)}}\ \exp\left[-\tfrac{1}{2}\left(\frac{\sigma_\theta^{-2}a}{\sigma_\theta^{-2} + a}\right)(\theta_o - \mu_\theta)^2\right]$$

$$> [100 - \mathrm{PR}(z_o)] - \frac{[100 - \mathrm{PR}(z_2)]}{\sqrt{\sigma_\theta^2(\sigma_\theta^{-2} + b)}}\ \exp\left[-\tfrac{1}{2}\left(\frac{\sigma_\theta^{-2}b}{\sigma_\theta^{-2} + b}\right)(\theta_o - \mu_\theta)^2\right],"$$

(4.14)

where

$$z_o = \frac{\theta_o - \mu_\theta}{\sqrt{\sigma_\theta^2}}$$

$$z_1 = \frac{\theta_o - \dfrac{a\theta_o + \sigma_\theta^{-2}\mu_\theta}{a + \sigma_\theta^{-2}}}{(a + \sigma_\theta^{-2})^{-\frac{1}{2}}} = \frac{\sigma_\theta^{-2}}{(a + \sigma_\theta^{-2})^{\frac{1}{2}}}\,(\theta_o - \mu_\theta)$$

$$z_2 = \frac{\theta_o - \dfrac{b\theta_o + \sigma_\theta^{-2}\mu_\theta}{b + \sigma_\theta^{-2}}}{(b + \sigma_\theta^{-2})^{-\frac{1}{2}}} = \frac{\sigma_\theta^{-2}}{(b + \sigma_\theta^{-2})^{\frac{1}{2}}}\,(\theta_o - \mu_\theta)\ .$$

We prefer d_2 when the direction of the inequality is reversed. Although Equation (4.14) looks rather frightening, it is really rather simple to use once you know the parameters of the utility function and of the posterior distribution on θ . Of course, those decision makers with access to a computer will find its application trivial.

5. A Three Action Example

Section 4 was built around the analysis of a decision problem in which two actions were available to the decision maker. Several families of utility functions were studied and the ideas were illustrated with an example involving the decision to advance or retain a student at a certain level in a sequential chain of instructional steps. Although the notation becomes more complex and the computation a bit more tedious, there are no fundamentally new ideas when we assume that there are three (or any finite number) of options open to the decision maker. In this section, we will illustrate this somewhat more general problem by using natural extensions of two of the families of utility functions discussed earlier.

5.1 Threshold Utility

Consider the following slight modification of the Individually Prescribed Instruction example discussed in Section 2.1. In the previous example, when a student completed a unit of instruction, he was considered a master or a nonmaster and was advanced or retained on the basis of expected utility. In this example, we merely extend the number of levels of mastery by further partitioning the nonmasters into two groups. The first group contains those nonmasters whose ability is close to the cutoff point separating the masters from the nonmasters. The second group contains those who apparently missed the whole point of the lesson. The state of a student being a nonmaster of the poorer variety will be denoted by θ_1; the better nonmasters will be denoted by θ_2; and, the masters by θ_3 .

For purposes of this example, we assume that there are only three actions available to the decision maker. The student may repeat both the present and the previous instructional units; he may repeat only the present unit; or, he may advance to the next unit.

With these specifications, we may now define the usual utility function by the following table.

Table 21

$u(d_i, \theta_j)$

	θ_1	θ_2	θ_3
d_1(back one)	u_{11}	u_{12}	u_{13}
d_2(retain)	u_{21}	u_{22}	u_{23}
d_3(advance)	u_{31}	u_{32}	u_{33}

As we have mentioned before, determination of these utilities is not an easy matter. In Section 3, we described one paradigm for their determination which might be helpful. However, we do not claim that it is the last word in utility specification. Nevertheless, in what follows, we will assume that the decision maker has coherently specified the utilities.

After the test score x is available, the decision will be made by selecting that action d_i, i = 1, 2, 3 which maximizes the posterior expected utility

$$\sum_{j=1}^{3} u(d_i, \theta_j)p(\theta_j|x) \quad .$$

We might think of this problem in terms of specifying two cutting test scores x_o and x_1, where $x_o < x_1$. Then for $x < x_o$, action d_1 will be taken; for $x_o < x < x_1$, action d_2 will be taken; and, for $x > x_1$, action d_3 will be taken. To determine the critical points x_o and x_1 which will divide the range of test scores into a d_1, a d_2, and a d_3 region, we return to a technique described in Section 2. We consider the posterior distribution of θ, $p(\theta|x)$, as a

function of x . Since x_o is the indifference point with respect to decisions d_1 and d_2, at x_o the expected posterior utility under d_1 must equal the expected posterior utility under d_2 . That is,

$$\sum_{j=1}^{3} u_{1j}p(\theta_j|x_o) = \sum_{j=1}^{3} u_{2j}p(\theta_j|x_o) \ .$$

Simplifying this, we see that x_o should be determined so that

$$(u_{11} - u_{13} + u_{23} - u_{21})p(\theta_1|x_o) + (u_{12} - u_{22} - u_{13} + u_{23})p(\theta_2|x_o) + (u_{13} - u_{23}) = 0 \ . \qquad (5.1)$$

Similarly, x_1 should be determined so that

$$(u_{31} - u_{33} + u_{23} - u_{21})p(\theta_1|x_1) + (u_{32} - u_{22} - u_{33} + u_{23})p(\theta_2|x_1) + (u_{33} - u_{23}) = 0 \ . \qquad (5.2)$$

In order to illustrate how to use Equations (5.1) and (5.2) in applications, we return to the example in Section 2.3 where posterior to our observation x, the ability parameter θ was continuous and, in fact, normally distributed. There we described the posterior distribution of θ by

$$p(\theta|x) \sim N\left[\frac{\tau\sigma^2 + x\phi}{\sigma^2 + \phi} \ , \ \frac{\phi\sigma^2}{\sigma^2 + \phi}\right] \ .$$

For purposes of this example, we also redefine the mastery levels θ_1, θ_2, and θ_3 in terms of critical points T_1 and T_2 on the ability scale (θ). We let $\theta_1 = \{\theta|\theta < T_1\}$, $\theta_2 = \{\theta|T_1 < \theta < T_2\}$, and $\theta_3 = \{\theta|\theta > T_2\}$. Then transforming the posterior $p(\theta|x)$ into a posterior on the normal deviate z, we see that

$$\Pr(\theta_1|x) = \Pr[z < z(T_1, x)],$$

$$\Pr(\theta_2|x) = \Pr[z < z(T_2, x)] - \Pr[z < (T_1, x)],$$

and

$$\Pr(\theta_3|x) = 1 - \Pr[z < z(T_2, x)]$$

where

$$z(T_i, x) = \frac{T_i - \frac{\tau\sigma^2 + x\phi}{\sigma^2 + \phi}}{[\phi\sigma^2/(\phi + \sigma^2)]^{\frac{1}{2}}} .$$

So we may rewrite Equations (5.1) and (5.2) in the form

$$(u_{11} + u_{22} - u_{21} - u_{12})\Pr[z < z(T_1, x_o)]$$

$$+ (u_{12} - u_{22} - u_{13} + u_{23})\Pr[z < z(T_2, x_o)]$$

$$+ (u_{13} - u_{23}) = 0$$

and

$$(u_{31} + u_{22} - u_{21} - u_{32})\Pr[z < z(T_1, x_1)]$$

$$+ (u_{32} - u_{22} - u_{33} + u_{23})\Pr[z < z(T_2, x_1)]$$

$$+(u_{33} - u_{23}) = 0 \quad .$$

Each of these equations now needs to be solved iteratively for x_o and x_1. It is recommended that T_1 and T_2 be used as first approximations to x_o and x_1, respectively.

We now turn to a modification of an example considered in Section 2.3 to illustrate these ideas. Assume that we have the following situation:

(1) $\rho(\theta) \sim N(80, 25)$

(2) $p(x|\theta) \sim N(\theta, 16)$

(3) $T_1 = 60$ and $T_2 = 85$

(4)

$u(d_i, \theta_j)$

	$\theta < T_1$	$T_1 < \theta < T_2$	$\theta > T_2$
d_1 (back one)	7	4	0
d_2 (retain)	2	6	1
d_3 (advance)	1	3	5

Thus, by applying Bayes theorem, the posterior distribution of θ as a function of x may be written in the form:

$$p(\theta|x) \sim N\left[\frac{80(16) + 25x}{16 + 25}, \frac{25(16)}{25 + 16}\right] .$$

Substituting into our equations for $z(T_1, x)$, we have

$$z(T_1, x) = \frac{60 - \frac{1280 + 25x}{41}}{3.123} = 9.214 - .195x$$

and

$$z(T_2, x) = \frac{85 - \frac{1280 + 25x}{41}}{3.123} = 17.218 - .195x .$$

And we must solve the equations

$$.07\ PR(9.214 - .195x_o) - .01\ PR(17.218 - .195x_o) - 1 = 0$$

$$.02\ PR(9.214 - .195x_1) - .07\ PR(17.218 - .195x_1) + 4 = 0 ,$$

where PR(z) equals the percentile rank of z . Iterating to a solution, we find that $x_o = 52.7$ and $x_1 = 87.4$. So the decision maker will choose decision d_1 and have the student repeat two units if $x \leq 52$; will choose decision d_1 and have the student repeat the current unit if $53 \leq x \leq 87$; and, will choose decision d_3 and advance the student if $x \geq 88$.

5.2 Linear Utility

Analogous to the situation in the two-action problem (see Section 4.2), we define linear utility to be linear in θ for each decision separately. Thus, linear utility in the three decision situation is defined by a function of the form:

$$u(d_i, \theta) = \begin{cases} e + f\theta & i = 1 \\ g + h\theta & i = 2 \\ k + m\theta & i = 3 \end{cases} . \qquad (5.3)$$

If we assume that our decisions can be indexed so that decision d_1 is most desirable when θ is small, so that decision d_2 is most desirable when θ takes intermediate values, and so that decision d_3 is most desirable when θ is large; then the solution of the three action problem is a straightforward extension of that offered in Section 4.2. Applying our permissible positive linear transformation, for $m > 0$, we let $u^*(d_i, \theta) = [u(d_i, \theta) - k]/m$. Thus,

$$u^*(d_i, \theta) = \begin{cases} e' + f'\theta & i = 1 \\ g' + h'\theta & i = 2 \\ \theta & i = 3 \end{cases}$$

where

$$e' = \frac{e - k}{m} \; ; \; f' = \frac{f}{m} \; ; \; g' = \frac{g - k}{m} \; ; \text{ and } h' = \frac{h}{m} \; .$$

And since we have four constants to estimate in specifying our linear utility function, we need four ordered pairs (θ_i, θ_j) such that $u^*(d_i, \theta_i) = u^*(d_j, \theta_j)$. Two of these pairs are provided by the breakeven points T_1 and T_2. At these points, we have

$$e' + f'T_1 = g' + h'T_1$$

and

$$g' + h'T_2 = T_2 .$$

Thus, we need only two additional pairs to completely specify the utility function. The resulting linear system of four equations in four unknowns can then be solved for e', f', g', and h'.

When we turn to maximizing expected utility, we now have three equations to consider. In fact, depending upon whether $e' + f'\mu_\theta$, $g' + h'\mu_\theta$, or μ_θ is largest, we choose decision d_1, d_2, or d_3, respectively. Graphically, this is clearly illustrated in Figure 5.1. All the decision maker needs to do is examine the utility of each decision at the mean of the posterior distribution of θ, choosing that decision with the highest value.

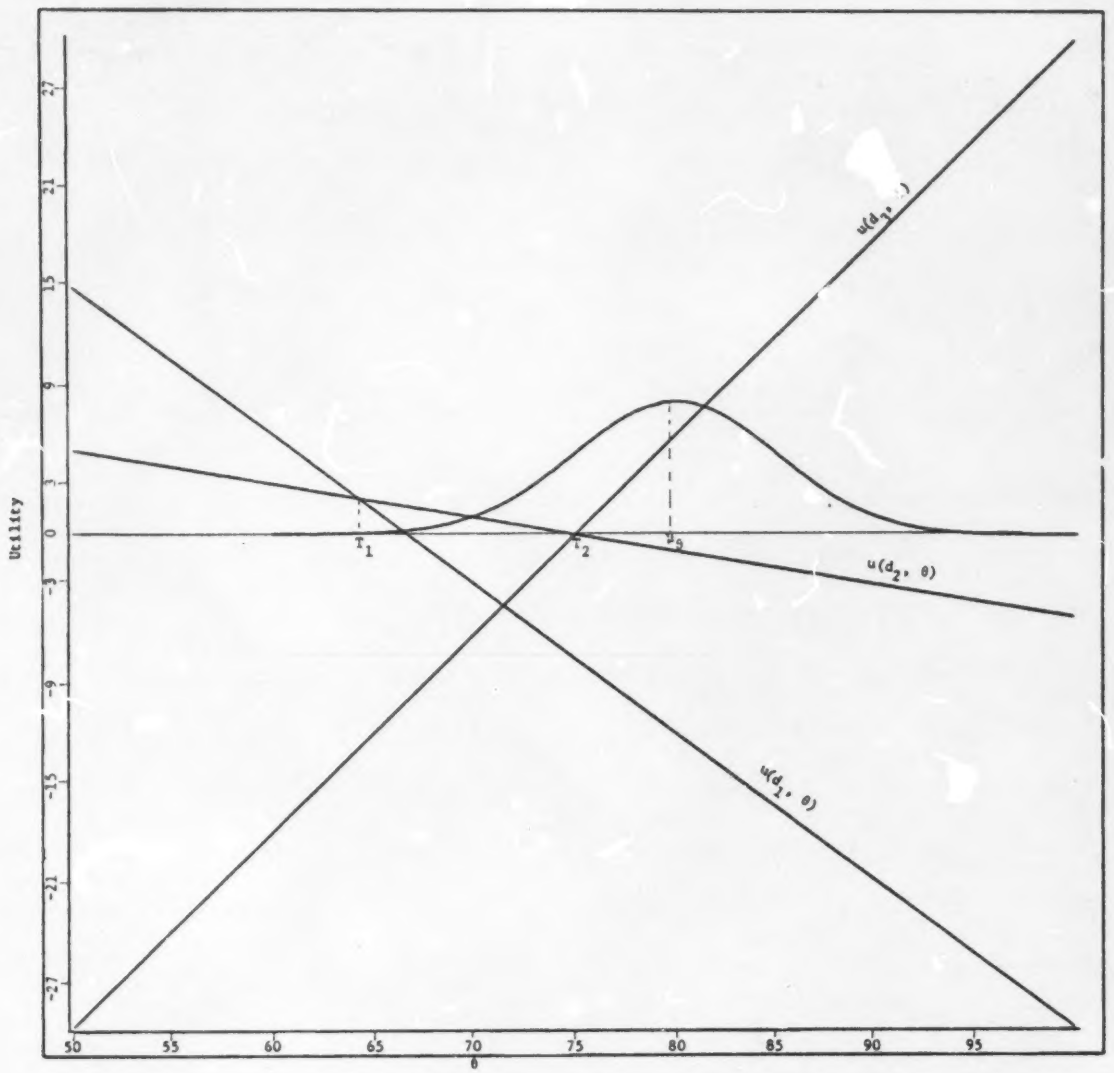

Figure 5.1

Linear Utility in a Three-Action Problem

This figure illustrates linear utility in the form of Equation (5.3) with constants e = 60, f = -.9, g = 15, h = -.2, k = -.90, and m = 1.2. As in the two action problem discussed in Section 4.2, the decision depends only upon the utilities of each decision at the posterior mean μ_θ. In this illustration, since $u(d_3, \mu_\theta) > u(d_2, \mu_\theta) > u(d_1, \mu_\theta)$, extensive form analysis will lead our decision maker to choose action d_3. In terms of T_1 and T_2, action d_1 will be taken whenever $\mu_\theta < T_1$; action d_2 will be taken whenever $T_1 < \mu_\theta < T_2$; and, action d_3 will be taken whenever $\mu_\theta > T_2$.

6. Preposterior Analysis

Information is never free. For example, information about the mastery level attained by a particular student is obtained by testing, interviews, or class recitation. Such activities spend the time of the student, spend the time of the teacher, and tie up facilities. If there are only meager rewards and penalties for correct and incorrect decisions, it may be wasteful to purchase information whose cost may exceed the gain in expected utility.

Suppose that a decision maker has (1) a prior distribution on θ, $\rho(\theta)$, (2) a utility function $u(d_i, \theta)$, and (3) a potential experiment which, if carried out, will have outcomes x with model density $p(x|\theta)$. Before collecting the data, the decision maker wants to know the extent to which his efforts are likely to be rewarded. That is, he wants to know whether the additional information contained in the potential experiment is likely to be sufficiently "valuable" to justify obtaining it. Bayesian decision analysis provides the framework of preposterior analysis for studying this question.

The logic of preposterior analysis is simple and can be readily understood by considering the following outline:

(1) The decision maker can attach a "value" to the information contained in his prior, $\rho(\theta)$, by calculating the expected utility of the optimal decision. That is,

$$\text{Value } [\rho(\theta)] = \max_{d_i} \int u(d_i, \theta)\rho(\theta)d\theta \ .$$

(2) Assume for the moment that the experiment has already been carried out and the result x obtained. If this were the case, then analogous to the above, the decision maker could attach a value to the information contained in his posterior. That is,

$$\text{Value } [p(\theta|x)] = \max_{d_i} \int u(d_i, \theta)p(\theta|x)dx \ .$$

(3) Continuing as if the data had already been collected, our decision maker could now calculate the "value added" by the experimental results (i.e., the increase in expected utility after the addition of the data).

$$\text{Value Added} = \max_{d_i} \int u(d_i, \theta)p(\theta|x)d\theta \; - \max_{d_i} \int u(d_i, \theta)\rho(\theta)d\theta \quad .$$

(4) Since, in fact, the experiment has not as yet been executed, of course, the value added cannot be determined. Nevertheless, the decision maker may consider the value added to be a function of the observation random variable X. In the jargon of decision theory, value added considered as a function of the random variable X is referred to as the <u>conditional value of sample (experimental) information</u> and is denoted:

$$v(e, x) = \max_{d_i} \int u(d_i, \theta)p(\theta|x)d\theta - \max_{d_i} \int u(d_i, \theta)\rho(\theta)d\theta \; .$$

It is conditional because it can be calculated only when x is known.

(5) Now since X is a random variable with a probability distribution $p(x) = \int \rho(\theta)p(x|\theta)d\theta$, it is clear that v(e, x) is also a random variable. If the density of v(e, x) were a simple function, it would be useful at this point to examine its location parameters and even credibility intervals. These statistics would describe the decision maker's prior beliefs about the probable increases in utility to be gained from sampling. In most applications, however, the density of v(e, x) is not a simple function. Although this complexity precludes most descriptive indices, in many instances, it will be possible to determine the mean of the distribution of v(e, x). In the decision theory literature, this

mean is commonly referred to as the <u>expected value of sample information</u> and is denoted

$$v(e) = \int v(e, x)p(x)dx$$

where $p(x)$ is the marginal prior density of X.

(6) The decision maker may now compare the expected value of sample information with the "cost" of performing the experiment and judge whether or not the experiment is likely to be worthwhile.

Before illustrating preposterior analysis with a numerical example, one central point must be made. In step (5), the decision maker must compare an expected utility with the cost of obtaining experimental information. It is critical that these two quantities not only be measured in the same units, but also that their respective scales have the same origin. If the expected value of experimental information, $v(e)$, is measured in arbitrary "utile" units while the cost of that information is in dollars and cents, a sensible comparison cannot be expected.

We will illustrate preposterior analysis with an example. We let

$$u(d_i, \theta) = \begin{cases} -(3/5)(\theta - 75) & i = 1 \\ (7/5)(\theta - 75) & i = 2 \end{cases} .$$

As in our previous examples, 75 has been selected as the indifference point between the acts of retaining (d_1) and advancing (d_2) the student. As expected, the advance decision (d_2) is positively related to ability (θ) while the retain decision (d_1) has a negative relationship.

Suppose further that the prior information about θ has been quantified in the form of a normal distribution with mean 78 and variance 36. Recall that in Section 4.2, we demonstrated that with linear utility, the optimum

decision depends only upon whether or not the mean of θ is greater than θ_o. Obviously in this case, with only prior information at the decision maker's disposal, the advance action will result in the highest expected utility. Thus,

$$\text{Value}[p(\theta)] = \int (7/5)(\theta - 75)p(\theta)d\theta$$

$$= u(d_2, \mu_{\theta(\text{prior})}) = (7/5)(\mu_{\theta(\text{prior})} - 75)$$

where $\mu_{\theta(\text{prior})}$ is the mean of the prior on θ. If the experimenter had carried out the experiment, the highest expected utility using the experimental results to help select the action would be

$$\text{Value } [p(\theta|x)] = \begin{cases} -\int (3/5)(\theta - 75)p(\theta|x)d\theta & \text{if } \mu_{\theta(\text{post})} < 75 \\ \int (7/5)(\theta - 75)p(\theta|x)d\theta & \text{if } \mu_{\theta(\text{post})} \geq 75 . \end{cases}$$

Since the utility function is linear in θ, the expected utility of decision d_i is merely the utility of the expectation or mean of θ. Thus,

$$\text{Value } [p(\theta|x)] = \begin{cases} u(d_1, \mu_{\theta(\text{post})}) = -(3/5)(\mu_{\theta(\text{post})} - 75) \text{ if } \mu_{\theta(\text{post})} < \\ u(d_2, \mu_{\theta(\text{post})}) = (7/5)(\mu_{\theta(\text{post})} - 75) \text{ if } \mu_{\theta(\text{post})} \geq 7 \end{cases}$$

The conditional value of sample information may be given by

$$v(e, x) = \begin{cases} u(d_1, \mu_{\theta(\text{post})}) - u(d_2, \mu_{\theta(\text{prior})}) \text{ if } \mu_{\theta(\text{post})} < 75 \\ u(d_2, \mu_{\theta(\text{post})}) - u(d_2, \mu_{\theta(\text{prior})}) \text{ if } \mu_{\theta(\text{post})} \geq 75 . \end{cases}$$

And the expected value of sample information is given by

$$v(e) = \int_{\substack{\text{all } x \\ \text{such that} \\ \mu_{\theta(post)} < 75}} [u(d_1, \mu_{\theta(post)})]p(x)dx + \int_{\substack{\text{all } x \\ \text{such that} \\ \mu_{\theta(post)} \geq 75}} [u(d_2, \mu_{\theta(post)})]p(x)dx - \int u(d_2, \mu_{\theta(prior)})p(x)dx \quad .$$

Thus, in order to evaluate v(e), all our decision maker needs to do is partition the range of x into two subsets: The first containing all x which will lead to a posterior mean which is less than 75, and the second containing those x which force $\mu_{\theta(post)}$ to be greater than 75. If, for the purposes of this example, we assume that the likelihood of our sample may be described by a normal distribution with mean θ and variance 25, then applying Bayes theorem, we see that the posterior distribution of θ is of the form

$$p(\theta|x) \sim N\left[\frac{36\ x + 25*78}{25 + 36}\ ,\ \frac{36*25}{36 + 25}\right] \quad .$$

And so the relationship $\mu_{\theta(post)} < 75$ is equivalent to the relationship $x < 72.92$. Thus, the expected value of sample information is

$$v(e) = \int_{-\infty}^{72.92} u(d_1, \mu_{\theta(post)})p(x)dx + \int_{72.92}^{+\infty} u(d_2, \mu_{\theta(post)})p(x)dx - u(d_2, \mu_{\theta(prior)}) \int_{-\infty}^{+\infty} p(x)dx \quad .$$

Let us pause here for just a moment and examine this equation. The first thing to notice is that we are integrating over a range of test scores X from $-\infty$ to $+\infty$. Conceptually, this may seem a little troublesome, for in most applications, test scores are bounded within a relatively small range.

Recall, however, that we assumed the model density was normal in form. This assumption implies that every x (from $-\infty$ to $+\infty$) has positive probability. Therefore, each x must be considered when taking the expectation. Although admittedly, this is a problem conceptually, in applications it is not very important. As long as the prior distribution on θ is carefully specified, there should be effectively zero probability that x will lie outside its permissible range.

In this particular example, we have $p(\theta) \sim N(78, 36)$ and $p(x|\theta) \sim N(\theta, 25)$. Therefore, $p(x) \sim N(78, 61)$. So that in this case, X has very little probability of falling outside the range (55, 101). Returning to our expression for the expected value of sample information, we find two integrals of the form:

$$\int u(d, \mu_{\theta(post)})p(x)dx \ .$$

We know that $u(d_i, \mu_{\theta(post)})$ is linear in $\mu_{\theta(post)}$. Since $\mu_{\theta(post)}$ is linear in x, this implies that $u(d_i, \mu_{\theta(post)})$ is also linear in x . In fact, by substituting $\mu_{\theta(post)} = (36\ x + 25 * 78)/(36 + 25)$, $u(d_i, \mu_{\theta(post)})$ may be written in the linear form

$$u(d_i, \mu_{\theta(post)}) = \begin{cases} -(3/5)(.59x - 43.03) & \text{if } i = 1 \\ (7/5)(.59x - 43.03) & \text{if } i = 2 \ , \end{cases}$$

and so, v(e) can be written as

$$v(e) = \int_{72.92}^{+\infty} (7/5)(.59x-43.03)p(x)dx - \int_{-\infty}^{72.92} (3/5)(.59x-43.03)p(x)dx - u(d_2, \mu_{\theta(prior)})$$

where $p(x) \sim N[78, 61]$. It can be shown that the following relationship holds.

$$\int_{-\infty}^{a} b(cx - d)p(x)dx = \frac{b(c\mu - d)}{100} PR(a) + b\sigma n\left[\frac{\mu - a}{\sigma} ; 0, 1\right]$$

where

$$p(x) \sim N[\mu, \sigma^2]$$

$PR(z)$ = percentile rank of z

$n\left[\frac{\mu - a}{\sigma} ; 0, 1\right]$ = the ordinate or height of the unit normal curve at $\frac{\mu - a}{\sigma}$.

Using this relationship, we may rewrite v(e) once again

$$v(e) = (7/5)\frac{(.59 \cdot 78 - 43.03)}{100}[100 - PR(72.92)] + (7/5)\cdot\sqrt{61}\, n[-.65: \ 0, 1]$$

$$-(3/5)\frac{(.59 \cdot 78 - 43.03)}{100} \quad PR(72.92) - (3/5)\cdot\sqrt{61}\, n[-.65: \ 0, 1]$$

$$-(7/5)(78 - 75)$$

$$= .46 \ .$$

References

Cronbach, L. J., & Gleser, G. C. Psychological Tests and Personnel Decisions. (1st Edition). Urbana: University of Illinois Press, 1957. (2nd Edition, 1965).

De Groot, M. H. Optimal Statistical Decisions. New York: McGraw-Hill, 1970.

Edwards, W., Lindman, H., & Savage, L. J. Bayesian statistical inference for psychological research. Psychological Review, 1963, 70, 193-242.

Hambleton, R., & Novick, M. R. Toward an integration of theory and method for criterion-referenced tests. Research Report No. 54. Iowa City, Iowa: The American College Testing Program, 1972.

Lewis, C., Wang, M., & Novick, M. R. Marginal distributions for the estimation of proportions in m groups. ACT Technical Bulletin No. 13. Iowa City, Iowa: The American College Testing Program, 1973.

Lindley, D. V. Bayesian statistics: A review. CMBS, Regional Conference Series in Applied Mathematics, 1972, No. 2.

Novick, M. R. Computer assisted data analysis. ACT Technical Bulletin No. 3. Iowa City, Iowa: The American College Testing Program, 1972.

Novick, M. R., Lewis, C., & Jackson, P. H. The estimation of proportions in m groups. Psychometrika, 1973, 38, 19-46.

Novick, M. R., & Jackson, P. H. Statistical Methods for Educational and Psychological Research. New York: McGraw-Hill, 1974.

Raiffa, H., & Schlaifer, R. Applied Statistical Decision Theory. Boston: Division of Research, Harvard Business School, 1961.

Wang, M. Tables of constants for the posterior marginal estimates of proportions in m groups. ACT Technical Bulletin No. 14. Iowa City, Iowa: The American College Testing Program, 1973.